



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)


(51) International Patent Classification ⁶ : H04L 12/46, 12/44	A1	(11) International Publication Number: WO 99/48252 (43) International Publication Date: 23 September 1999 (23.09.99)
(21) International Application Number: PCT/US98/15554 (22) International Filing Date: 24 July 1998 (24.07.98) (30) Priority Data: 08/907,385 7 August 1997 (07.08.97) US (71) Applicant: VIXEL CORPORATION [US/US]; Suite 100, 15245 Alton Parkway, Irvine, CA 92618-2307 (US). (72) Inventor: BERMAN, Stuart, B.; 2010 Vista Caudal, Newport Beach, CA 92660 (US). (74) Agents: MURPHY, David, B. et al.; Lyon & Lyon LLP, Suite 4700, 633 West Fifth Street, Los Angeles, CA 90071-2066 (US).		(81) Designated States: CA, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>With amended claims and statement.</i>

(54) Title: METHODS AND APPARATUS FOR FIBRE CHANNEL INTERCONNECTION OF PRIVATE LOOP DEVICES

(57) Abstract

Methods and apparatus for Fibre Channel interconnection is provided between a plurality of private loop devices through a Fibre Channel private loop device interconnect system. In the preferred embodiments, the Fibre Channel private loop device interconnect system is a fabric or an intelligent bridging hub. In one aspect of this invention, a Fibre Channel private loop device is connected to two or more Arbitrated Loops containing, or adapted to contain, one or more private loop devices. Preferably, the interconnect system includes a routing filter to filter incoming Arbitrated Loop physical addresses (ALPAs) to determine which Fibre Channel frames

must attempt to be routed through the fabric. Numerous topologies of interconnect systems may be achieved. In another aspect of this invention, a method is provided for implementing a logical loop of private loop devices by segmenting the logical loop into a plurality of sets, assigning each set to a physical Arbitrated Loop and connecting the Arbitrated Loops to a Fibre Channel private loop device interconnect system. Additional methods are provided for restricting attached devices to Arbitrated Loop physical addresses within certain ranges. Additionally, methods are provided for resetting hosts, the method generally comprising the steps of detecting the addition of a storage device to a first Arbitrated Loop, and thereafter, resetting the Arbitrated Loop or loops on which a host or hosts reside on second Arbitrated Loop. Methods for operation with use of SCSI initiators generate a link service reject when no address match is found, or when an address match is found, but where no device with the destination ALPA exists on the Arbitrated Loop corresponding to the destination.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DESCRIPTIONMethods And Apparatus For Fibre Channel
Interconnection Of Private Loop Devices5 Field of the Invention

The present invention relates to input/output channel and networking systems, and more particularly to methods of using a Fibre Channel fabric or intelligent bridging hub to interconnect Fibre Channel Arbitrated Loops composed of private
10 loop devices (i.e., devices which do not support direct fabric attachment).

Related Application Information

This application is a continuation-in-part of Application Serial No. 08/801,471 filed February 18, 1997, entitled "Fibre
15 Channel Switching Fabric", now pending.

Background of the Invention

Fibre Channel is an American National Standards Institute (ANSI) set of standards which describes a high performance serial transmission protocol which supports higher level
20 storage and network protocols such as HIPPI, SCSI, IP, ATM and others. Fibre Channel was created to merge the advantages of channel technology with network technology to create a new I/O interface which meets the requirements of both channel and network users. Channel technology is usually implemented by
25 I/O systems in a closed, structured and predictable environment, whereas network technology usually refers to an open, unstructured and unpredictable environment.

Advantages of Fibre Channel typically include the following. First, it achieves high performance, which is a
30 critical in opening the bandwidth limitations of current computer to storage and computer to computer interfaces at speeds up to 1 gigabit per second or faster. Second, utilizing

fiber optic technology, Fibre Channel can overcome traditional I/O channel distance limitations and interconnect devices over distances of 6 miles at gigabit speeds. Third, it is high level protocol independent, enabling Fibre Channel to transport a wide variety of protocols over the same media. Fourth, Fibre Channel uses fiber optic technology which has a very low noise properties. Finally, cabling is simple in that Fibre Channel typically replaces bulky copper cables with small lightweight fiber optic cables.

10 Fibre Channel supports three different topologies: point-to-point, Arbitrated Loop and fabric attached. The point-to-point topology attaches two devices directly. The Arbitrated Loop topology attaches devices in a loop. The fabric attached topology attaches a device directly to a
15 fabric.

 The Arbitrated Loop topology was initially designed to provide a lower cost interconnect than fabrics and to provide more interconnect than point-to-point topologies. The Arbitrated Loop topology was created by separating the transmit and receive fibers associated with each loop port and connecting the transmit output of one loop port to the receive input of the next loop port. Typically, characteristics of the Arbitrated Loop topology include: first it, allows up to 126 participating node ports and one participating fabric port to
20 communicate, second, each node port implements a route filtering algorithm, and third, all ports on a single loop have the same upper 16 bits of the 24-bit NL_Port address identifier.

 There are two classifications of devices on an Arbitrated Loop: private loop devices and public loop devices. Public loop devices attempt a Fabric Login (FLOGI) upon initialization. Public loop devices also are cognizant of all twenty four bits of the 24-bit NL_Port native port address identifier. Public loop devices will open the fabric port at
30 Arbitrated Loop Physical Address (ALPA, bits 7 to 0) zero when
35

the domain and area (bits 23 to 8) do not match their domain and area. Private loop devices use only the lower eight bits of the ALPA and can only communicate within the local loop.

Generally, the disadvantages of the Arbitrated Loop topology include: first, it is a blocking topology, that is, only a single connection between a pair of nodes is allowed at any point in time (excluding the broadcast mode). Second, device buffering occurs in each device as it has a six word buffer, creating a delay of up to 225 nanoseconds. This delay is additive with each device in the loop. The delay creates overhead for the communicating devices when a large number of devices are connected to a loop. Third, distance also adds delay to a loop and is additive for each device. For copper medium there is a 4 nanosecond delay per meter and for optical medium there is a 5 nanosecond delay per meter. Fourth, robustness is an issue since all devices are on one loop any device failure will cause the entire loop to fail or reset. Fifth, the total bandwidth available is limited to the bandwidth of the loop itself. Finally, device failure is an issue since while frames are being transmitted, a timeout in an upper level protocol may occur, thereby disrupting the applications.

Loop devices are typically interconnected on an Arbitrated Loop with a hub, see FIG. 22 numeral 678. The hub is a passive device, that is a loop exists within the hub 674, 675, 676, 677, 679. A hub in most cases maintains the loop's integrity when devices are removed, powered off, or fail by using a port bypass circuit 674, 675, 676, 677. Hubs simply receive and redrive the signals to individual devices.

There are many disadvantages which result when interconnecting private loop devices with hubs: First, hubs do not address the blocking nature of the loop topology. Second, jitter is propagated from bypassed nodes. This additive affect causes loop instability when a large number of devices are interconnected. Third, when data is currently

being transferred and a device attached to a hub is powered off or fails, the loop could be reset which is destructive to the communicating devices. Fourth, if a device is inserted into a live loop the loop will be reset which is destructive to the communicating devices.

The majority of initial Fibre Channel equipment deployment utilizes the Arbitrated Loop topology with hubs as the interconnect. These environments are experiencing all the previously defined problems inherent in both Arbitrated Loop topology and with hub deployment. The blocking nature of the Arbitrated Loop is limiting the number of devices on a loop.

The distance and delay parameters are also creating more overhead for the loop. Finally the loop is being reset by single devices.

As such, it is the goal of this invention to provide apparatus and methods which solve or mitigate these problems.

Summary of the Invention

This invention relates to methods and apparatus for Fibre Channel interconnection of a plurality of private loop devices through a Fibre Channel private loop device interconnect system. In the preferred embodiments, the Fibre Channel private loop device interconnect system is a fabric or an intelligent bridging hub. Through these methods and apparatus, multiple Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, may be interconnected even though on separate Arbitrated Loops.

In the preferred embodiment, an interconnection system is provided for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops either containing, or being adapted to contain, one or more private loop devices. The apparatus preferably includes at least a first Arbitrated Loop containing, or adapted to contain, one or more private loop devices, and at least a second Arbitrated Loop, either containing, or adapted to contain, one or more private loop

devices. The Arbitrated Loops are interconnected via a Fibre Channel private loop device interconnect system which is disposed between the Arbitrated Loops.

5 In one embodiment, a Fibre Channel fabric is disposed between the Arbitrated Loops, and includes a routing filter which filters incoming Arbitrated Loop physical addresses (ALPAs) to determine which Fibre Channel frames must attempt to be routed through the fabric. Thus, by disposing routing information regarding private loop devices within the Fibre
10 Channel private loop device interconnect system, multiple Arbitrated Loops containing private loop devices may be interconnected.

Any type of private loop device, consistent with the apparatus and methods stated herein, may be utilized in
15 conjunction with this system. Examples of private loop devices include storage devices, such as tape drives, JBODs and RAID subsystems, host systems, and other connections within a system, such as bridges, particularly SCSI to Fibre Channel bridges, routers, particularly Fibre Channel to asynchronous
20 transfer mode systems and Fibre Channel to ethernet systems.

Various interconnection topologies may be utilized with these systems. Beyond a single fabric having two Arbitrated Loops, any number of Arbitrated Loops may be utilized, consistent with the size constraints of the fabric.
25 Alternately, a first fabric, with at least one Arbitrated Loop attached, and a second fabric, with at least one Arbitrated Loop, may have the first fabric and second fabric directly connected. Alternatively, or additionally, a first fabric and a third fabric, each having at least one Arbitrated Loop, may
30 be connected through a second fabric. Yet another topology includes a first fabric having M ports, where one port is connected to storage, preferably JBODs, and the remaining M-1 ports of the first fabric are connected to M-1 second fabrics.

An alternative interconnect topology includes a first fabric
35 with M ports, and M second fabrics, each second fabric being

connected to the first fabric. In yet another interconnect topology, a device is connected to a first fabric by a first path and to a second fabric by a second, independent path.

5 In yet another interconnect topology, a first set of fabrics and a second set of fabrics may be interconnected through a first intermediate fabric and a second intermediate fabric, each of the first sets of fabrics connected to the first intermediate fabric, and separately to the second intermediate fabric, and each of the second set of fabrics connected to the
10 first intermediate fabric and separately to the second intermediate fabric.

In yet another aspect of this invention, a method is provided for implementing a logical loop of private loop devices in a novel manner. The method generally comprises the
15 steps of segmenting the logical loop of private loop devices into a plurality of sets, assigning each set to a physical Arbitrated Loop and connecting the Arbitrated Loops to a Fibre Channel private loop device interconnect system to effect interconnection of the Arbitrated Loops.

20 In another aspect of this invention, a method is provided for selectively filtering Fibre Channel frames. This method serves to route frames between one or more private loop devices on a first Arbitrated Loop and one or more private loop devices on at least a second Arbitrated Loop. Preferably, the method
25 includes the steps of receiving the Fibre Channel frames over the first Arbitrated Loop at a connected port of a Fibre Channel private loop device interconnect system and filtering the frame by, either, forwarding the frame on the first Arbitrated Loop if the frame has an address on the first
30 Arbitrated Loop, or, providing an "open" response on the first Arbitrated Loop if the address is not on the first Arbitrated Loop. Optionally, in the event that the frame includes an address not on the first Arbitrated Loop, the additional step of attempting to route the frame through the Fibre Channel
35 private loop device interconnect system may be made. In yet

another optional step, buffering of the frames destined to private loop devices not on the first Arbitrated Loop may be performed, most preferably, permitting cut-through if the route can be made without substantial buffering.

5 Yet another novel method of these inventions is a method for restricting attached devices to Arbitrated Loop physical addresses (ALPAs) within certain ranges. Through this method, multiple Fibre Channel Arbitrated Loops of private loop devices are configured, each private loop device on the Arbitrated Loop
10 having an Arbitrated Loop physical address. Generally, the steps in the preferred method comprise, first, dividing the ALPAs into nonoverlapping sets, second, assigning each set to a separate physical Arbitrated Loop, and thereafter, during loop initialization, forcing the attached private loop devices
15 to choose from the assigned set.

Yet another novel method comprises a method for resetting hosts within a Fibre Channel interconnection system of private loop devices. In this method of operation of an interconnection system, the system including more than one
20 Arbitrated Loop, at least one loop being adapted to contain storage and one loop adapted to contain a host, those devices being private loop devices, the loops being connected to a Fibre Channel private loop device interconnect system, the method generally comprises the steps of, first, detecting at
25 least the addition of a storage device to a first Arbitrated Loop, and thereafter, resetting the Arbitrated Loop or loops on which a host or hosts reside upon such detected addition.

In this manner, a host resident on an Arbitrated Loop becomes aware of storage private loop devices which have been added to
30 other Arbitrated Loops separated from the host bearing Arbitrated Loop by a Fibre Channel private loop device interconnect system.

In yet another method of operation of the inventive system, a method for operation with use of SCSI initiators is
35 provided. In this interconnection system, the system includes

more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop devices, the loops being connected to a Fibre Channel private loop device interconnect system. The method generally comprises the steps of first, receiving port login (PLOGI) input/output (I/O) probes at the Fibre Channel private loop device interconnect system, thereafter, performing address look-up for the received PLOGI I/O probes, and, if a match exists in the look up, routing the PLOGI I/O probes from the Fibre Channel SCSI initiator to private loop devices on the Fibre Channel private loop device interconnect system or other Fibre Channel private loop device interconnect system. In the event that no match is found upon address look up, the PLOGI I/O probes are routed to the Fibre Channel private loop device interconnect system controller, and a link service reject (LS_RJT) is returned.

Similarly, a link service reject is returned in the event that an address match is found, but where no device with the destination ALPA exists on the Arbitrated Loop corresponding to the destination.

An intelligent bridging hub adapted to interconnect a plurality of Arbitrated Loops containing private loop devices is provided. The intelligent bridging hub includes at least first and second hub submodules, the submodules comprising a plurality of ports, the ports including port bypass circuits connected to the ports for connecting to the Arbitrated Loops adapted to contain private loop devices, and, an Arbitrated Loop physical address filtering port, a router, the router being disposed between the first and second hub submodules, and a processor control coupled to the router and the first and second submodules. The router need not support all classes of Fibre Channel connections, for example, the router may optionally not support class 1 connections. Optionally, the processor control need not provide back-up route determination mechanisms.

Objects of the Invention

Accordingly, it is an object of this invention to interconnect separate Arbitrated Loops of private loop devices through Fibre Channel private loop device interconnect system.

5 It is yet a further object of this invention to segment one logical loop composed of private loop devices into several physical Arbitrated Loops each of which is connected to Fibre Channel private loop device interconnect system.

10 It is yet a further object of this invention to connect private loop devices over a fabric without any a priori knowledge by those devices or special software driver modifications to support the Fibre Channel private loop device interconnect system.

15 It is yet a further object of this invention to route the I/O probes from a Fibre Channel SCSI Initiator to private loop devices on other ports on the Fibre Channel private loop device interconnect system or on other ports on connected Fibre Channel private loop device interconnect systems.

20 It is yet a further object of this invention to handle I/O probes from Fibre Channel SCSI Initiators which are destined for nonexistent devices.

25 It is yet a further object of this invention to filter the frames received by a Fibre Channel private loop device interconnect system port and select those frames which are destined to private loop devices on separate fabric ports or for other ports on connected Fibre Channel private loop device interconnect system.

30 It is yet a further object of this invention to reset private loop hosts when storage devices are added or removed on other ports on the Fibre Channel private loop device interconnect system or on other ports on connected Fibre Channel private loop device interconnect system.

It is yet a further object of this invention to limit the ALPA range that connected private loop devices can choose when

connected to a Fibre Channel private loop device interconnect system port.

Brief Description of the Drawings

FIG. 1 is a block diagram illustrating the use of the Stealth Mode in connecting one logical loop of private loop devices by segmenting it into several physical loops all interconnected through a network of fabrics.

FIG. 2 is a block diagram illustrating the bandwidth advantage in connecting multiple pairs of simultaneously communicating private loop devices.

FIG. 3 is a block diagram showing the additive delays of interconnecting several devices in a loop.

FIG. 4 is a block diagram showing the Stealth Mode procedure for routing Fibre Channel SCSI Initiator PLOGI I/O probes which do not correspond to an existent ALPA.

FIG. 5 is a block diagram showing the fabric router modifications necessary to route Fibre Channel SCSI Initiator PLOGI I/O probes which do not correspond to an existent ALPA.

FIG. 6 is a block diagram showing the interconnection of seven host computers with ten racks of JBODs, each JBOD consisting of up to eight disk drives.

FIG. 7 is a block diagram showing the interconnection of thirteen host computers with ten racks of JBODS, each JBOD consisting of up to eight disk drives.

FIG. 8 is a block diagram showing the interconnection of forty nine host computers with five racks of JBODS, each JBOD consisting of up to eight disk drives.

FIG. 9 is a block diagram showing the interconnection of a JBOD and a RAID storage subsystem with four hosts. Each device is connected through a redundant path.

FIG. 10 is a block diagram showing the interconnection of forty eight hosts with eight RAID storage subsystems.

FIG. 11 is a block diagram showing the interconnection of four hosts with a JBOD and a RAID storage subsystem. There are redundant links to the JBOD and the RAID subsystems.

FIG. 12 is a block diagram showing the interconnection of
5 forty hosts with eight RAID storage subsystems. There are redundant paths for all connections which span at least two fabrics.

FIG. 13 is a block diagram illustrating the use of a Fibre Channel Fabric.

10 FIG. 14 is a block diagram of a Fibre Channel Fabric.

FIG. 15 is a block diagram of the Fabric Control module.

FIG. 16 is a block diagram of the fabric Router.

FIG. 17 is a block diagram of the fabric Port Control.

FIG. 18 is a diagram of the LIFA frame format.

15 FIG. 19 is a diagram of an example LIFA frame to restrict ALPA selection to the range 1x, i.e., 10, 17, 18, 1B, 1D, 1E and 1F.

FIG. 20 is a diagram of a LISM frame with a World-wide name of all zeros.

20 FIG. 21 is a block diagram of a fabric with eight ports, showing the port locations.

FIG. 22 is a block diagram of a four port Fibre Channel hub.

25 FIG. 23 is a block diagram of a JBOD (Just a Bunch of Disks) storage subsystem.

FIG. 24 is a diagram illustrating the relationship between a logical Arbitrated Loop and a physical Arbitrated Loop. The physical Arbitrated Loop is segmented over a fabric or an intelligent bridging hub.

30 FIG. 25 is a block diagram of an intelligent bridging hub.

Detailed Description of the InventionTable of Contents

	A.	Definitions
	B.	Stealth Mode Features and Functions
5	C.	Background - Fibre Channel Fabric Model
	1.	Fabric Control Module
	2.	Fabric Router
	a.	Address Table
	b.	Address Match Module
10	c.	Route Determination Module
	3.	Switch Core
	4.	Port Control Module
	D.	Background - Intelligent Bridging Hub
	E.	ALPA Filtering
15	F.	Port ALPA Range Configuration
	G.	Fibre Channel SCSI Initiator I/O Probe Spoofing
	H.	Resetting Host Ports Upon Storage Addition
	I.	Default Route Configuration
	J.	Stealth Loop Topologies and Configuration Strategies
20	K.	Other Documents

A. Definitions

For expository convenience, the present invention in various aspects is referred to as the Stealth Loop Mode, the lexicon being devoid of a succinct descriptive name for a system of the type hereinafter described. The following discussions will be made clearer by a brief review of the relevant terminology as it is typically (but not exclusively) used.

The "Fibre Channel ANSI standard" describes the physical interface, transmission protocol and signaling protocol of a high-performance serial link for support of the higher level protocols associated with HIPPI, IPI, SCSI, IP, ATM and others.

The Fibre Channel Fabric comprises hardware and software that switches Fibre Channel frames between attached devices at speeds up to one gigabit per second.

5 "FC-1" defines the Fibre Channel transmission protocol which includes the serial encoding, decoding, and error control.

"FC-2" defines the signaling protocol which includes the frame structure and byte sequences.

10 "FC-3" defines a set of services which are common across multiple ports of a node.

"FC-4" is the highest level in the Fibre Channel standards set. It defines the mapping between the lower levels of the Fibre Channel and the IPI and SCSI command sets, the HIPPI data framing, IP, and other Upper Level Protocols (ULPs).

15 A "fabric" (sometimes referred to as a switch or router) is an entity which interconnects various N_Ports attached to it and is capable of routing frames by using the Destination Identifier (D_ID) information in the FC-2 frame header.

20 An "intelligent bridging hub" is a hub with one or more ports which implement ALPA filtering and/or routing functions and contains control logic.

A "FC-PLD-IS" or Fibre Channel Private Loop Device Interconnect System is a fabric or intelligent bridging hub, or a device which achieves the functionality of these devices.

25 A "RAID" or redundant array of inexpensive disks storage device is an interleaved storage technique which speeds access to disks along with implementing redundant storage access methods.

30 A "JBOD" or Just A Bunch Of Disks is a storage subsystem composed of a series of disks. A JBOD is similar to a RAID system without the RAID controller which implements the RAID striping and mirroring features.

"Topology" is an interconnection scheme that allows multiple Fibre Channel ports to communicate. For example

point-to-point, Arbitrated Loop and fabric-attached are all Fibre Channel topologies.

"Fabric topology" is a topology where a device is directly attached to a fabric and that uses the Destination Identifier (D_ID) embedded in the frame header to route the frame through a Fabric to the desired destination N_Port.

"Point-to-point topology" allows communication between two N_Ports without the use of a Fabric.

"Arbitrated Loop topology" permits two or more L_Ports to communicate using arbitration to establish a point-to-point circuit. When two L_Ports are communicating, the Arbitrated Loop topology supports simultaneous, symmetrical bidirectional data flow.

"Stealth Mode" is a fabric or intelligent bridging hub mode of operation which allows the interconnection of private loop devices over multiple fabric/hub ports.

"Port" is a generic reference to an N_Port or F_Port.

"Link Control Facility" is a facility which attaches to an end of a link and manages transmission and reception of data. It is contained within each Port type.

An "N_Port" is a hardware entity which includes a Link Control Facility.

An "NL_Port" is an N_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology.

An "F_Port" is a generic reference to an F_Port or FL_Port.

An "FL_Port" is an F_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology.

An "L_Port" is an N_Port or F_Port that contains Arbitrated Loop functions associated with Arbitrated Loop topology.

A "Node" is a collection of one or more N_Ports controlled by a level above FC-2.

A "frame" is an indivisible unit of information used by FC-2.

"Classes of Service" are different types of services provided by the Fabric and used by the communicating N_Ports.

"Class 1" service is a service which establishes a dedicated connection between communicating N_Ports.

5 "Class 2" service is a service which multiplexes frames at frame boundaries to or from one or more N_Ports with acknowledgment provided.

10 "Class 3" service is a service which multiplexes frames at frame boundaries to or from one or more N_Ports without acknowledgment.

A "Gigabit Link Module" is a module which interfaces to the Endec through either a 10-bit or 20-bit interface and interfaces to the Fibre Channel link through either a copper or fiber interface.

15 An "Encoder/Decoder" or Endec is a device which implements the FC-1 layer protocol.

"Fabric Login Protocol" is when an N_Port interchanges Service Parameters with the Fabric by explicitly performing the Fabric Login protocol or implicitly through an equivalent
20 method not defined in FC-PH.

"Private Loop Device" is a device which does not attempt a fabric login (FLOGI) ELS command and cannot open a fabric port, e.g., ALPA zero, when the domain and area addresses of a frame to be transmitted is not equal to zero.

25 "Public Loop Device" is a device which attempts fabric login and can communicate with devices that contain nonzero domain and area address values. Public loop devices can observe the rules of either public or private loop behavior.

30 A public loop device may communicate with both private and public loop devices.

"Private Loop Direct Attach" or PLDA is a technical report which defines a subset of the relevant standards suitable for the operation of peripheral devices such as disks on a private loop.

"N_Port Login" or PLOGI is a Fibre Channel Extended Link Service Command defined in the FCPH Revision 4.3 ANSI standard that requests transfer of Service Parameters from the initiating N_Port/NL_Port to the N_Port/NL_Port or
5 F_Port/FL_Port associated with the Destination Identifier.

"World-wide Name" or WWN is an 8-byte field which uniquely identifies an N_Port or F_Port. Each N_Port or F_Port must have a WWN which is unique worldwide.

"N_Port Identifier" is a 3-byte native address field
10 which is unique within the Fibre Channel address domain.

"Source Identifier" or S_ID is the address identifier used to indicate the source Port of the transmitted frame.

"Destination Identifier" or D_ID is the address identifier used to indicate the targeted destination of the
15 transmitted frame.

"Link Services Reject" or LS_RJT is a Fibre Channel Extended Link Service Command defined in the FCPH Revision 4.3 ANSI standard that notifies the transmitter of a Link Service request that the Link Service request Sequence has been
20 rejected. LS_RJT frames may be transmitted for a variety of conditions which may be unique to a specific Link Service Request.

"Loop Initialization" is a protocol used to initialize the loop prior to beginning loop operations or when
25 configuration changes are detected.

"Loop Initialization Primitive" or LIP frames are used to accomplish loop initialization.

"Loop Initialization Select Master" or LISM frame contains the devices World-wide unique name (WWN). The device
30 with the lowest WWN becomes the temporary loop master during the loop initialization sequence.

"Loop Initialization Fabric Address" or LIFA loop primitive allows public ports that had logged-in with the fabric to reclaim the ALPA they had been using immediately
35 prior to loop initialization. The LIFA is the first loop

primitive transmitted by the temporary loop master which contains ALPA information.

5 "Loop Initialization Report Position" or LIRP loop primitive is used to build a map of all loop devices relative to the loop master. This primitive is optional in loop initialization.

10 "Loop Initialization Loop Position" or LILP loop primitive allows any interested device on the loop to view the current loop map. This primitive is optional in loop initialization.

"Jitter" is random variation in the timing of a signal, especially a clock.

B. Stealth Mode Features & Functions

15 The Stealth Mode is a combination of features and functions applied to a fabric or an intelligent bridging hub which allows it to interconnect private loop devices between ports on the fabric or hub, preferably without the private loop device's knowledge. The result of this mode allows one logical Arbitrated Loop to be segmented into multiple physical loops.

20 As shown in FIG. 24 a logical loop 709 of Arbitrated Loop devices 700, 701, 702, 703 is connected physically 705, 706, 707, 708 to a fabric or intelligent bridging hub 704.

This segmentation solves some or all of the problems inherent in Arbitrated Loop topologies, such as:

- 25 • Multiple loops allow more than two devices to communicate simultaneously, thereby increasing the bandwidth of the entire logical loop. As shown in FIG. 2, Host number 41 and JBOD 43 can simultaneously communicate 45 while Host number 42 and RAID 44 communicates 46.
- 30 • When a private loop device powers up or down or resets, only the port in which that device is connected is affected (except for certain conditions, see Resetting Host Ports Upon Storage

Addition section, below). In FIG. 4, if the Host connected to port 1 numeral 70 powers down, only that port is affected. The loop state of the devices connected to ports 2, 3 and 4 are all unaffected.

- Jitter is not additive for the entire logical loop, just the physical loop directly attached to the fabric port.
- Device buffering delay is minimized since there are less devices per physical loop. As shown in FIG. 3 each device on the loop adds six words of buffering and associated delay 51, 52, 53 and 54.
- Delay introduced from interdevice distance is minimized since the delay will only affect the devices on the single fabric port. As shown in FIG. 3 the distance between devices 59, 60, 61, 62 creates an additive delay.
- Existing Fibre Channel SCSI Initiator and IP (Internet protocol) software drivers and Fibre Channel native disk drive firmware do not have to be modified.

Some or all of the following features are implemented by the fabric or intelligent bridging hub to create the Stealth Mode.

- The fabric port looks to the loop devices as a device with a large amount of ALPAs. The fabric accomplishes this by applying a routing filter to the incoming ALPAs.
- The fabric predetermines the ALPA range available for each port and forces the connected devices to choose that ALPA range during loop initialization to simplify system routing.
- Upon initialization the Fibre Channel SCSI Initiators send out PLOGI I/O probes to all ALPAs to find any attached SCSI devices. The fabric routes

the I/O probes to actual attached SCSI devices and redirects those I/O probes to nonexistent ALPAs to the fabric controller where they are responded to.

- The fabric may reset host ports when storage is added to a fabric port to initiate Fibre Channel SCSI Initiator PLOGI I/O probing.
- Block transmission of the LIRP and LILP frames during loop initialization in case host driver software modifies its PLOGI I/O probing based on the ALPA map information contained in those frames.

While a majority of the description in this application describes the implementation via a Fibre Channel fabric, because that is our preferred embodiment, the inventions are equally applicable to Intelligent Bridging hubs or any Fibre Channel private loop device interconnect system..

C. Background - Fibre Channel Fabric Model

As shown in Fig. 13, a Fibre Channel Fabric is an entity which transports Fibre Channel frames between attached devices. The data transmission between the connected device port (i.e., N_Port) and the Fabric port (i.e., F_Port) is serial and consists of one or more frames. The transmission protocol and speeds along with the fabric functionality are defined in the American National Standard for Information Systems (ANSI) FCPH standard (see Other Documents section, below).

A primary function of the Fabric is to receive frames from a source N_Port and route the frames to the destination N_Port whose address identifier is specified in the frames. Each N_Port is physically attached through a link to the Fabric or in the case of an Arbitrated Loop topology attached to the same loop. FC-2 specifies the protocol between the Fabric and the attached N_Ports. A Fabric is characterized by a single address space in which every N_Port has a unique N_Port identifier.

The Fabric model contains three or more F_Port or FL_Ports. Each F_Port is attached to an N_Port through a link. Each F_Port is bidirectional and supports one or more communication models. The receiving F_Port responds to the sending N_Port according to the FC-2 protocol. The Fabric may or may not verify the validity of the frame as it passes through the Fabric. The Fabric routes the frame to the F_Port directly attached to the destination N_Port based on the N_Port identifier (D_ID) embedded in the frame. The address translation and the routing mechanisms within the Fabric are transparent to N_Ports.

FIG. 13 shows a possible environment containing a Fibre Channel fabric. The fabric is identified by the reference numerals 401 and 402. The fabrics illustrated are connected with a mix of workstations 403, disk arrays 404, mainframe computers 405, and Personal Computers (PC) 406. Fabric interconnection is not limited to particular equipment or a network topology as illustrated in FIG. 13. Two types of fabric topologies are illustrated in FIG. 13., the direct fabric attached topology 409 and the Arbitrated Loop topology 407. The fabrics in FIG. 13 are shown interconnected or networked through a link 408. All links to the fabric can operate at either 266 Mbps, 533 Mbps or 1.063 Gbps speeds and operate over either copper or fiber media.

FIG. 14 shows a block diagram of the fabric. The fabric is composed of a fabric control module 454, a router module 452, multiple port control modules 451, 474, 475 a switch core module 453 and optionally one or more brouter modules 455. The Fabric Control module controls and configures the rest of the fabric but is not usually involved in the normal routing of frames. The fabric Router 452 performs route address matching, route determination based on the ANSI X3T11 rules, route request blocking & unblocking, switch core programming 463, statistics collection and port control module route request/response handling 459, 460, 461, 462, 466, 467, 472,

473. The fabric Port Control modules (PCM) 451, 470, 474, 475 receive Fibre Channel frames from the fiber or copper media 456, 477, 478, perform frame validation, send a route request to the router 459, 461, 466, 472, receives a route response from the router 460, 462, 463, 467, 473, forwards the frame to the switch core 457, 469, and either discards the frame, modifies the frame into a fabric reject (F_RJT) or fabric busy (F_BSY) frame or forwards the frame depending on the route response from the router. The fabric switch core 453 is a nonblocking NxN matrix switch with 36 bit wide transmit and receive I/Os. The switch core switches frames from the PCMs 451, 470, 474, 475 to the destination PCMs or Brouter Module.

1. Fabric Control Module

FIG. 14 shows the Fabric Control module (FCM) 454. The FCM configures the fabric, collects and reports network management parameters and implements the fabric defined servers such as the Simple Name Server, Directory Services, etc. The FCM configures the router 452, the port control modules 451, 474, 475 and the brouter module 455. FIG. 15 shows the Fabric Control module (FCM) in more detail. The FCM is made up of fast SRAM 482, DRAM 483, a DUART 484, flash memory 485 (nonvolatile storage), a processor 481 and a Decode/DMA Control module 487. The code for the processor is contained in the flash memory 485 and is copied to SRAM upon bootup. The interface to the brouter module 455 allows the FCM to communicate through legacy networks such as ethernet and fast ethernet, depending on the brouter module.

2. Fabric Router

The Fabric Router, FIG. 14 numeral 452 receives route requests generated from the Port Control modules 459, 461, 466, 472, determines the frame route, reports the route responses to the Port Control modules 460, 462, 467, 473, programs the switch core to connect and disconnect the routes 463, manages

blocked route requests and collects the routing statistics. There is one central router contained in a fabric. The Router connects and disconnects routes on a frame by frame basis. Since the router can determine a route in real time (i.e.,
5 Fibre Channel frame time) the Fabric can support Class 1 frames. The router is realized in hardware through either an FPGA or a custom ASIC. The router is composed of thirteen functional modules as illustrated in FIG. 16:

- Port Control Route Request Interface (PCRRIM) 530
- 10 • Port Control Route Response Interface (PCRSPM) 544
- Address Table 532
- Address Match Module (ADM) 531
- Blocked Route Request Table Module (BRTBL) 533
- Blocked Route Request Port Register Array (BRRP) 534
- 15 • Blocked Route Request Timer (BRTMR) 535
- Route Request Unblock Determination Module (RRUNB) 536
- Route Request Selector (RRS) 537
- Route Determination Module (RDM) 538
- Route State Table (RST) 539
- 20 • Router Statistics Gathering Module (RST) 541
- Router Control FSM (RCFSM) 540.

a. Address Table

The Address Table is shown in FIG. 16 numeral 532. The address table is initially configured by the processor in the
25 fabric control module 522. The Address Table contains entries against which the incoming Fibre Channel frame destination identifier (D_ID) is compared. The address entry contains a twenty four bit address mask register along with a twenty four
30 bit address register. The incoming D_ID is ANDed with the address mask register and the result is compared to the address register. This allows a match to be performed on any number of bits in the address. This also implements routing based on any combination of the address domain (upper eight bits of the address field), area (middle eight bits of the address field)

or port (lower eight bits of the address field) fields. Additional address fields include the destination port and the address priority fields. The destination port indicates which remote F_Port to route the frame to and the address priority field specifies a priority for this address table entry match. For any two address matches the address table entry match which is the highest priority will be used. This implements alternate routing in case of port failure feature.

b. Address Match Module

The Address Match module (ADM) is shown in FIG. 16 numeral 531. The ADM performs the comparison with the incoming frame D_ID address from the route request 505 with the Address Table contents 509. The results are used by the Route determination module 538. The ADM has as an input the twenty-four bit address to match, i.e., the incoming frame D_ID address from the route request, and returns the following responses: the remote match port, the address matched indication and the route to control module indication. The ADM will match an incoming D_ID address to all the addresses in the address table in one clock. The ADM logic is implemented in combinatorial logic. The ADM performs the following checks for each address table entry:

$$\text{Address Match indication} = (\text{address in table} == (\text{address mask} \& \text{D_ID}))$$

The results are then priority decoded based on address priority contained in the address table and the resulting address match signal and port are generated. There is one special mode which is implemented which will preemptively route all frames to the Fabric Control module except frames originating from the Fabric Control module. This allows the fabric control module to process all incoming frames which is useful when the fabric is functioning in certain environments.

c. Route Determination Module

FIG 16 numeral 538 shows the Route Determination module (RDM). The RDM applies rules defined in the ANSI Fibre Channel specifications to calculate how to route the incoming frame.

5 The RDM receives the route request 510 from the RRS 537 along with route context for the source and destination ports 512 from the Route State Table 539. The RRS outputs the route results 545, 511 to both the Router Control FSM 540 and the PCRSPM 544. The RDM is implemented in combinatorial logic and
10 applied the route rules in one clock.

3. Switch Core

FIG. 14 shows the Switch Core. The switch core implements a nonblocking NxN matrix switch. The input to the switch core comes from the individual Port Control modules FIG.

15 14 numerals 457 and 469. The output from the switch core is wired to the Endec FIG. 14 numeral 458 and the Brouter Module FIG. 14 numeral 476. The switch core is paths are setup and torn down by the router FIG. 14 numeral 463.

4. Port Control Module

20 FIG. 14 shows the Port Control (PC) locations 451, 470, 474, 475, within the fabric block diagram. There is one PC per port or link. The PC interfaces with the fabric attached device through either copper or fiber media 456, 477, 478. The PC interfaces to the switch core through transmit 458 and receive
25 457 data buses and control signals. The PC interfaces to the router through route request 459, 461, 466, 472 and route response 460, 462, 467, 473 buses and control signals. Finally the PC interfaces to the Fabric Control module through a processor interface bus 465.

30 FIG. 17 shows the Port Control in more detail. Frames are received from the fiber or copper link 551 and enter the Endec 553. The Endec implements the 8B/10B encoding/decoding,

the loop port state machine and fabric/point-to-point state machine functions and outputs thirty two bit data words with two bits of parity and tag information to the receive FIFO 555. The PC contains a module which guards against a receive FIFO
5 overrun 154 condition. Once the receive FIFO 555 starts filling, the Port Control Module (PCM) 556 reads the frame header, requests a route from the router 563, 564 and forwards the frame to the switch core 561, 562. The PCM is configurable by the processor 570 in the Fabric Control module. The Port
10 Control also receives frames from the switch core 565, 566 to be transmitted by the Endec 553.

D. Background - Intelligent Bridging Hub

An intelligent bridging hub is a device composed of one or more passive hubs interconnected by some additional logic
15 to bridge between two or more Fibre Channel Arbitrated Loops.

An intelligent bridging hub can implement the Stealth routing mode if it contains the following functionality:

- Route filtering for two or more hub submodules
- Basic routing between hub submodules
- 20 • Minimum processor functions to "spoof" the PLOGI I/O probes and participate in loop initialization.

A block diagram of an intelligent bridging hub is shown in FIG. 25. As shown in FIG. 25, the intelligent bridging hub 773 is composed of two or more hub submodules 774, 775
25 containing a route filtering port 759, 760, some logic to perform limited routing 761 and a processor 762 to perform loop initialization and some other stealth features. Each hub submodule contains port bypass circuits 751 through 758 or their equivalent.

30 E. ALPA Filtering

To receive frames for attached physical Arbitrated Loop segments the fabric port implements receive frame ALPA range filtering. This filtering function is done in the

encoder/decoder module, see FIG. 17 numeral 553 of the port control logic, FIG. 14, numerals 451, 474, 475. As each frame is received, the fabric port applies a mask to the received ALPA (Arbitrated Loop Physical Address) and compares it with
5 a preconfigured value. Depending on the "receive on match/no match" bit the fabric will receive the frame on a resultant match or no match. This algorithm is shown below.

```
if (((received frame ALPA & mask) == predetermined  
address) == (match/no match))
```

```
10     receive frame
```

```
    else
```

```
        forward frame to next device
```

The receive frame on match/no match bit is used to allow greater filtering flexibility of incoming frames. An example
15 of this filtering is shown by the following example. Assume there is one private loop device attached to the fabric port and its ALPA is 17h (where h means hexadecimal notation). To route all frames from this device to the fabric the fabric mask would be 00, the predetermined address would be 00h and the
20 port would receive frame "on match". This has the affect of filtering, i.e., receiving, all frames transmitted from the attached device.

Another example would be a hub connected to a fabric port with attached hub device ALPAs of B1h, B2h, B3h, B4h, B5h,
25 B6h, B9h, BAh, BCh. The fabric mask would be F0, the predetermined address would be B0h and the port would receive frame on "no match". The result would have the fabric port receiving all frames that do not contain B in bits 7 to 4 of the ALPA.

30 There can be multiple ALPA filters per port. An example would be a port with a fabric mask of F0 and a predetermined address of both 20h and 40h. If the mask algorithm returns a positive result when applied to any of multiple filters for a single port (i.e., an OR result) the frame is forwarded on the
35 local loop and not filtered.

While the preceding description is of the preferred embodiment, you can still achieve some of the benefits of the invention without the capability of a settable match/no match bit. Even if you do not support a match/no match bit you can
5 still support numerous topologies of interconnected Arbitrated Loops composed of private loop devices with a fabric or intelligent bridging hub.

F. Port ALPA Range Configuration

In order to support the Stealth Mode the fabric must have
10 a priori knowledge of the device ALPA ranges on each port. This is accomplished by the fabric by forcing the attached devices to choose a predetermined ALPA range. This is done during the LIFA loop initialization phase. The fabric reserves all ALPAs in the LIFA bit map, see FIG. 18 numeral 600, by
15 setting them equal to 1, except the range that the fabric desires the port to choose from. FIG 19 shows the ALPA bit map 601 to exclude all ALPAs from being chosen from the attached devices except 10, 17, 18, 1B, 1D, 1E or 1F.

The fabric is guaranteed to generated the LIFA by
20 becoming loop master in the LISM phase. The fabric does this by choosing the lowest World-wide name, i.e., zero in the LISM frame, see FIG. 20 numeral 610.

G. Fibre Channel SCSI Initiator I/O Probe Spoofing

Private loop Fibre Channel SCSI Initiator devices send
25 out N_Port Logins (PLOGI) to the entire ALPA range after loop initialization to probe for SCSI devices. The PLOGIs are transmitted serially. After each PLOGI transmission the Fibre Channel SCSI Initiator waits for each reply before sending another PLOGI. If the host driver receives the PLOGI it has
30 just sent that indicates there are no devices on the loop with the ALPA.

Since a fabric in the Stealth Mode is filtering and routing frames off the local Arbitrated Loop, in many cases the

Fibre Channel SCSI Initiator will not receive its own PLOGIs destined for nonexistent ALPAs and will timeout before sending another one. Since this timeout can be up to ten seconds and there are 126 possible devices on a loop the initialization time is not acceptable unless the fabric acts on the PLOGI frame.

In the Stealth Mode the fabric is optimized to automatically route the PLOGIs destined to nonexistent ALPAs to the fabric controller which will return an immediate response. The ANSI FCPH standard requires all Class 3 frames which are not deliverable to be discarded. In the Stealth Mode the fabric deviates slightly from the ANSI standard and routes Class 3 frames to be discarded to the internal fabric controller.

If the frames received by the fabric controller are Class 3 PLOGI frames (i.e., Fibre Channel SCSI Initiator I/O probes) the fabric element will return a Link Services Reject (LS_RJT) to indicate that the exchange is not to be setup. All other Class 3 frames will be discarded by the fabric element satisfying the ANSI FCPH standard. This satisfies the requirement to return a frame to the PLOGI in real time to avoid PLOGI timeout.

Two modifications of the fabric router are necessary to handle PLOGI probes to nonexistent ALPAs. The first modification is to route all rejected Class 3 frames to the fabric controller. As shown in FIG. 4, all PLOGI's 78 which should be discarded by the router 76 due to the destination being nonexistent are routed to the fabric controller 75. The fabric controller generates an LS_RJT (Extended Link Services Reject) frame 80, 81 for every Class 3 PLOGI frames it receives. The LS_RJT frame is routed back to the originating port 82 to expedite the PLOGI probing phase.

The second modification is to route all Class 3 frames which were transmitted and received back from a loop (i.e., no device present) to the fabric controller. As shown in FIG. 5

if a PLOGI is generated from a loop 96 and routed to a destination loop 97, 98 which does not contain a device which matches the ALPA, it is received by the router 93 and routed to the fabric controller 95. As in the previous case the
5 fabric controller 95 generates an LS_RJT (Extended Link Services Reject) frame 101 for every Class 3 PLOGI frames it receives. That frame is routed back to the originating port 102, 103 to expedite the PLOGI probing phase.

H. Resetting Host Ports Upon Storage Addition

10 Fibre Channel SCSI Initiators only I/O probe with PLOGI frames whenever the loop is reset. If private loop storage devices are added to remote fabric ports in a Stealth Mode environment after private loop Fibre Channel SCSI Initiators are initialized, they will have missed the PLOGI I/O probe
15 phase and will not be "seen" by the SCSI Initiators. Fibre Channel SCSI Initiators must be notified of the addition of the storage device so as to reinitiate I/O probing with PLOGI frames. In the Stealth Mode the fabric implements an option to reset ports, i.e., transmit LIP, which have hosts attached
20 if a port with storage is added.

I. Default Route Configuration

The Stealth Mode requires the ALPAs for the devices attached to the fabric port to be predetermined in order to simplify fabric routing. The fabric enforces this
25 configuration by generating special LIFA frames during loop initialization, see Port ALPA Range Configuration Section. The strategy is to allocate numerically similar ALPA ranges to fabric ports to simplify routing. The ALPA ranges are chosen based on the number of ALPAs in certain ranges. ALPA ranges
30 are defined as 1x, 2x, 3x, 4x, 5x, 6x, 7x, 8x, 9x, Ax, Bx, Cx, Dx and Ex, where x is ALPA bits 3 to 0, and represents a don't care value. Note that all numbers are in hexadecimal notation. Since the ALPA values are not contiguous, ALPA range selection

must be done carefully. The table below organizes the ALPA values into ranges and shows the number of available ALPA addresses in each range.

ALPA Range	Number of Addresses in Range
00	1
01,02,04,08,0F	5
10,17,18,1B,1D,1E,1F	7
23,25,26,27,29,2A,2B,2C,2D,2E	10
31,32,33,34,35,36,39,3A,3C	8
43,45,46,47,49,4A,4B,4C,4D,4E	10
51,52,53,54,55,56,59,5A,5C	9
63,65,66,67,69,6A,6B,6C,6D,6E	10
71,72,73,74,75,76,79,7A,7C	9
80,81,82,84,88,8F	6
90,97,98,9B,9D,9E,9F	7
A3,A5,A6,A7,A9,AA,AB,AC,AD,AE	10
B1,B2,B3,B4,B5,B6,B9,BA,BC	9
C3,C5,C6,C7,C9,CA,CB,CC,CD,CE	10
D1,D2,D3,D4,D5,D6,D9,DA,DC	9
E0,E1,E2,E4,E8,EF	6

5

For example, JBOD (Just a Bunch of Disks) storage devices are typically composed of several Fibre Channel disk drives in a single 19" equipment rack mount enclosure, see FIG. 23. Each drive 680, 681, 682, 683 requires a separate ALPA. The ports which contain JBODs should use a range have enough ALPAs to assign to all drives. Hubs are another device which contain several addressable entities, see FIG. 22. Therefore only the ranges 2x, 3x, 4x, 5x, 6x, 7x, Ax, Bx, Cx and Dx can be assigned to ports which have JBODs, hubs or other fabrics attached. Hosts and RAID devices (Redundant Array's of Inexpensive Disks) only use one ALPA and can use any ALPA range. Fabric links use seven ALPAs and use ALPA ranges which contain seven or more ALPAs.

10

15

J. Stealth Loop Topologies and Configuration Strategies

To take full advantage of the Stealth Mode all private loop devices should be connected to the fabric in such a way to expedite simultaneous communication between pairs of devices. When multiple devices must share fabric ports, similar devices should share the same loops to simplify the routing requirements. For example storage devices should share the same loop and fabric port(s).

There are a large number of topologies which are supported by the Stealth Mode. In addition all types of private loop devices are supported such as hosts, RAIDs, JBODs, hubs, SCSI-to-Fibre Channel bridges, tape drives, other fabrics, etc. Below are descriptions of several representative examples of the Stealth Mode topologies.

In all diagrams the fabric shown contains eight ports. As shown in FIG. 21, port 1 641 is on the far left hand side and port 8 648 is on the far right hand side of the fabric. The example topologies are representative but not exhaustive. The examples show a single line between the fabric and the attached device(s). The single line represents an Arbitrated Loop connecting the fabric to the attached device(s), except for fabric-to-fabric links which are point-to-point and not Arbitrated Loop. In many cases the loop is only composed of two devices, the fabric port and the attached device. Although the examples show an eight port fabric, higher or lower port size fabrics or intelligent bridging hubs may also be used.

FIG. 1 shows a fabric interconnecting a mix of private loop devices including: JBOD systems 3, 6, 9, a hub 16, hosts, RAID systems 19, 29, 36, in addition to other fabrics 2, 37. Note that all the links between the fabric and attached devices are Arbitrated Loops, there connection is shown by a single line. The fabric routing table for the AGS/8 1 fabric is shown below. Note the primary/backup column indicates whether the

link should be used as the primary link or a backup link in case the primary link fails.

Port	Device Type	Address	Mask	Priority
1	JBOD, 2 racks 4, 5	Ax, Cx	F0	Primary
2	Hub 16	2x	F0	Primary
3	JBOD, 2 racks 7, 8	4x, 6x	F0	Primary
4	RAID 19	80	FF	Primary
5	Host 20	81	FF	Primary
6	JBOD, 5 racks 10, 11, 12, 13, 14	3x,5x,7x,Bx,Dx	F0	Primary
7	Fabric link 21	9x	F0	Primary
8	Fabric link 22	1x	F0	Primary

- 5 The port ALPA filtering table (i.e., of received frames) is shown below.

Port	Device Type	Address	Mask	Match
1	JBOD, 2 racks	Ax, Cx	F0	No Match
2	Hub	2x	F0	No Match
3	JBOD, 2 racks	4x, 6x	F0	No Match
4	RAID	80	FF	No Match
5	Host	81	FF	No Match
6	JBOD, 5 racks	3x,5x,7x,Bx,Dx	F0	No Match
7	Fabric link	9x	F0	No Match
8	Fabric link	1x	F0	No Match

FIG. 6 shows a fabric interconnecting a mix of private
 10 loop devices including: JBODs and host devices on a single fabric. The fabric routing table for the AGS/8 120 fabric is shown below.

Port	Host Name	Address	Mask	Filter on Match/No Match
1	Host 121	10	FF	Primary
2	Host 122	17	FF	Primary
3	Host 123	18	FF	Primary
4	Host 124	1B	FF	Primary
5	Host 136	1D	FF	Primary
6	Host 137	1E	FF	Primary
7	Host 138	1F	FF	Primary
8	JBOD, 10 racks 126,127,128,129, 130,131,132,133,134, 135	Cx,Ax,6x,4x,2x,Dx, Bx,7x,5x,3x	F0	Primary

The port ALPA filtering table is shown below.

Port	Host Name	Address	Mask	Filter on Match/No Match
1	Host	10	FF	No Match
2	Host	17	FF	No Match
3	Host	18	FF	No Match
4	Host	1B	FF	No Match
5	Host	1D	FF	No Match
6	Host	1E	FF	No Match
7	Host	1F	FF	No Match
8	JBOD, 10 racks	Cx,Ax,6x,4x,2x, Dx,Bx,7x,5x,3x	F0	No Match

5

FIG. 7 replaces host port number 7 in FIG. 6 with a fabric 141 containing seven more hosts 159, 160, 161, 162, 163, 164, 165. The routing table is similar to FIG. 6 except the routing address for port 7 is 9x, the routing mask is F0, the
 10 ALPA filtering address is 9x and the ALPA filtering mask is F0.

FIG. 8 replaces all hosts in FIG. 6 with fabrics 201, 202, 203, 204, 205, 206, 207 with attached host devices. The

ALPA ranges that the fabrics use are 1x, 9x, 3x, 5x, 7x, Bx and Dx. Because of limited ALPA ranges the maximum number of JBODs 208 in this configuration is five 209, 210, 211, 212, 213. The JBODs are assigned ALPAs of 2x, 4x, 6x, Ax and Cx.

5 FIG. 9 shows a topology utilizing two redundant fabrics 220, 241 to provide redundant links 229, 230, 231, 232, 233, 234, 235, 236, 237, 238, 239, 240 from hosts 223, 224, 225, 226 to either JBOD 221 or RAID 222 devices. Both the routing tables and the port ALPA filtering tables for the two fabrics 10 would be identical.

FIG. 10 shows a topology containing maximum single ALPA private loop devices interconnected with fabrics 250, 251, 252, 253, 254, 255, 256, 257, 258. There are no devices with multiple ALPA's in this topology such as JBODs or hubs.

15 FIG. 11 shows a topology of two storage systems, a JBOD 261 and a RAID 262, which are connected to a fabric with redundant links 269, 270 and 271, 272. FIG. 11 also shows interconnection of four host private loop devices 263, 264, 265, 266. Each storage system connection is composed of a 20 primary/active link 269, 271 and a backup link 270, 272.

FIG. 12 shows a topology similar to FIG. 10 except for redundant fabrics 300, 301 interconnecting the leaf or outer fabrics 302, 303, 304, 305, 306, 307, 308, 309. This topology allows redundant paths between each leaf fabric.

25 K. Other Documents

The following documents provide selected ANSI information regarding Fibre Channel technology:

- 1) ANSI X3.230-1994, "Fibre Channel Physical and Signaling Interface (FC-PH)".
- 30 2) ANSI X3.297-1996, "Fibre Channel Physical and Signaling Interface (FC-PH-2)".
- 3) ANSI X3.303-1996, "Fibre Channel Physical and Signaling Interface (FC-PH-3)".

4) ANSI X3.272-1996, "Fibre Channel Arbitrated Loop (FC-AL)".

5) ANSI X3T11 Project # 1162-DT, "Fibre Channel Private Loop Direct Attach (PLDA)".

5 6) ANSI X3T11 Project # 1133-D, "Fibre Channel Arbitrated Loop 2 (FC-AL-2)".

7) Kembel, R., "The Fibre Channel Consultant - Arbitrated Loop", Connectivity Solutions. ISBN 0-931836-82-4, 1996, 1997.

10 Although the foregoing invention has been described in some detail by way of illustration and example for purposes of clarity and understanding, it may be readily apparent to those of ordinary skill in the art in light of the teachings of this invention that certain changes and modifications may be made
15 thereto without departing from the spirit or scope of the appended claims.

Claims

1. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the Loops containing one or more private loop devices, comprising:

5 a first Arbitrated Loop containing one or more private loop devices,

one or more additional Arbitrated Loops containing one or more private loop devices, and

10 a Fibre Channel private loop device interconnect system disposed between the first Arbitrated Loop and an additional Arbitrated Loop.

2. An Interconnection system for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the Loops adapted to contain one or more private loop devices, comprising:

15 a first Arbitrated Loop adapted to contain one or more private loop devices,

one or more additional Arbitrated Loops adapted to contain one or more private loop devices, and

20 a Fibre Channel private loop device interconnect system disposed between the first Arbitrated Loop and an additional Arbitrated Loop.

3. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system is a Fiber Channel fabric.

4. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system is an intelligent bridging hub.

5. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system implements route filtering.

6. The interconnection system of claim 1 or 2 wherein
5 the Fibre Channel private loop device interconnect system implements port ALPA range configuration.

7. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system implements I/O probe spoofing.

10 8. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system resets a port which has a host connected to it upon the addition of storage on an Arbitrated Loop connected to an additional Fibre Channel private loop device interconnect
15 system port.

9. The interconnection system of claim 1 or 2 where there are 2 loops.

10. The interconnection system of claim 1 or 2 where there are more than 10 loops.

20 11. The interconnection system of claim 1 or 2 wherein the private loop devices include storage.

12. The interconnection system of claim 11 wherein the storage is JBODs.

13. The interconnection system of claim 11 wherein the
25 storage is RAIDs.

14. The interconnection system of claim 11 wherein the storage is tape drives.

15. The interconnection system of claim 1 or 2 wherein the private loop devices include hosts.

16. The interconnection system of claim 1 or 2 further including one or more public loop devices.

5 17. The interconnection system of claim 1 or 2 wherein the private loop devices includes a bridge.

18. The interconnection system of claim 17 wherein the bridge is a SCSI to Fibre Channel bridge.

10 19. The interconnection system of claim 1 or 2 wherein the private loop devices includes a router.

20. The interconnection system of claim 19 wherein the router connects Fibre Channel to asynchronous transfer mode (ATM).

15 21. The interconnection system of claim 19 wherein the router connects Fibre Channel to ethernet.

22. The interconnection system of claim 21 wherein the router connects Fibre Channel to fast ethernet.

23. The interconnection system of claim 21 wherein the router connects Fibre Channel to a gigabit ethernet.

20 24. The interconnection system of claim 1 or 2 wherein at least one arbitrated loop connection includes a hub.

25 25. The interconnection system of claims 1 or 2 wherein there are redundant connections between a Fibre Channel private loop device interconnect system and an attached storage unit or host unit.

26. An interconnection system for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

first Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop, and

a second Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop, the first Fibre Channel private loop device interconnect system and second Fibre Channel private loop device interconnect system being directly connected.

27. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop,

a second Fibre Channel private loop device interconnect system, and

a third Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop,

the first Fibre Channel private loop device interconnect system being directly connected to the second Fibre Channel private loop device interconnect system, which is directly connected to the third Fibre Channel private loop device interconnect system.

28. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first Fibre Channel private loop device interconnect system, with M ports,

M-1 second Fibre Channel private loop device interconnect system, the first Fibre Channel private loop device interconnect system directly connecting to each of the M-1 Fibre Channel private loop device interconnect system, and

a JBOD connected to the first Fibre Channel private loop device interconnect system.

29. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first Fibre Channel private loop device interconnect system with M ports, and

M second Fibre Channel private loop device interconnect system, each second Fibre Channel private loop device interconnect system connected to the first Fibre Channel private loop device interconnect system.

30. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first Fibre Channel private loop device interconnect system,

a second Fibre Channel private loop device interconnect system, and

at least a first device, the first device connected to the first Fibre Channel private loop device interconnect system and to the second Fibre Channel private loop device interconnect system.

31. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first set of Fibre Channel private loop device interconnect systems,

a second set of Fibre Channel private loop device interconnect systems,

a first intermediate Fibre Channel private loop device interconnect system, and

a second intermediate Fibre Channel private loop device interconnect system,

each of the first set of Fibre Channel private loop device interconnect systems connected to the first intermediate
5 Fibre Channel private loop device interconnect system, and separately to the second intermediate Fibre Channel private loop device interconnect system, and

each of the second set of Fibre Channel private loop device interconnect systems connected to the first intermediate
10 Fibre Channel private loop device interconnect system, and separately to the second intermediate Fibre Channel private loop device interconnect system.

32. A method for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing
15 one or more private loop devices, comprising the steps of:

providing a first Arbitrated Loop containing one or more private loop devices,

providing one or more Arbitrated Loops containing one or more private loop devices, and

20 disposing a Fibre Channel private loop device interconnect system between the first Arbitrated Loop and an additional Arbitrated Loop.

33. A method for implementing a logical loop of private loop devices, comprising the steps of:

25 segmenting the logical loop of private loop devices into a plurality of sets,

assigning each set to a physical Arbitrated Loop,

connecting the Arbitrated Loops to a Fibre Channel private loop device interconnect system to interconnect the
30 Arbitrated Loops.

34. A method for routing frames between a private loop device on a first Arbitrated Loop and one or more private loop

devices on at least a second Arbitrated Loop comprising the steps of:

receiving frames over the first Arbitrated Loop at a connected port of a Fiber Channel private loop device
5 interconnect system, and

filtering said frames by:

(1) forwarding the frame on the first Arbitrated Loop if it has an address on the first Arbitrated Loop and

(2) providing an "open" response on the first Arbitrated
10 Loop, if the address is not on the first Arbitrated Loop.

35. The method for routing frames of claim 34 further including the step of attempting to route the frame to the address which is not on the first Arbitrated Loop.

36. The method for routing frames of claim 34 further
15 including the step of buffering of the frames destined to private loop devices not on the first Arbitrated Loop if the route cannot be made immediately.

37. The method for routing frames of claim 34 further including the step of look at destination address and compare
20 with predefined table, then cut through if the route can be made immediately.

38. The method of claim 34 wherein the filtering is achieved by:

applying a mask to the received ALPA,
25 comparing the results with preconfigured value, and
comparing the result (match/no match) with a filter on the match/no match bit.

39. A method for configuring of multiple Fibre Channel Arbitrated Loops of private loop devices, each private loop

device on an Arbitrated Loop having an Arbitrated Loop Physical Address (ALPA), comprising the steps of:

- dividing the ALPAs into non-overlapping sets,
- assigning each set to a separate physical Arbitrated
- 5 Loop, and
- during loop initialization, forcing the attached private
- loop devices to choose from the assigned set.

40. The method of claim 39 further including the step of interconnecting each subplurality through a Fibre Channel

10 private loop device interconnect system.

41. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop

15 devices, the loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

- detecting an addition of a storage device to a first Arbitrated Loop, and
- resetting the Arbitrated Loops on which hosts reside upon
- 20 such detected addition.

42. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop

25 devices, the loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

- receiving PLOGI I/O probes at a Fibre Channel private loop device interconnect system,
- performing address lookup for the received PLOGI I/O
- 30 probes, and if a match exists in the lookup,
- routing I/O probes from Fibre Channel SCSI initiator to private loop devices on the Fibre Channel private loop device

interconnect system or other Fibre Channel private loop device interconnect systems.

43. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop devices, the loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

receiving PLOGI I/O probes at a Fibre Channel private loop device interconnect system,

performing address lookup for the received PLOGI I/O probes, and

routing the I/O probes to the Fibre Channel private loop device interconnect system controller and returning a link service reject if either there is no match found upon address lookup or if a match is found upon address lookup, but no device with the destination ALPA exists on the Arbitrated Loop corresponding to the destination.

44. The method of the preceding claim wherein the determination that no device with the destination ALPA exists on the Arbitrated Loop is performed by the steps of:

sending a class 3 frame on a loop, and

simultaneously receiving the same frame on the same loop.

45. A bridging hub for connection of a plurality of Arbitrated Loops, the Arbitrated Loops being adapted to contain private loop devices, comprising:

at least first and second hub submodules,

the submodules comprising:

a plurality of ports, the ports including port bypass circuits connected to the ports for connection to the Arbitrated Loops adapted to contain private loop devices, and an ALPA filtering port,

a router, the router disposed between the first and second hub submodules, and

a processor control coupled to the router and the first and second submodules.

5 46. The bridging hub of claim 45 wherein the router does not support class 1 connections.

47. The bridging hub of claim 45 wherein the processor control does not provide backup route determination mechanisms.

AMENDED CLAIMS

[received by the International Bureau on 08 February 1999 (08.02.99);
original claims 1-47 replaced by amended claims 1-47 (11 pages)]

1. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the Loops containing one or more private loop devices, comprising:
- 5 a first Arbitrated Loop containing one or more private loop devices,
one or more additional Arbitrated Loops containing one or more private loop devices, and
a Fibre Channel private loop device interconnect
10 system disposed between the first Arbitrated Loop and an additional Arbitrated Loop for routing frames between private loop devices on the first Arbitrated Loop and private loop devices on the additional Arbitrated Loop.
- 15 2. An Interconnection system for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the Loops adapted to contain one or more private loop devices, comprising:
- a first Arbitrated Loop adapted to contain one or more private loop devices,
20 one or more additional Arbitrated Loops adapted to contain one or more private loop devices, and
a Fibre Channel private loop device interconnect system disposed between the first Arbitrated Loop and an additional Arbitrated Loop for routing frames between
25 private loop devices on the first Arbitrated Loop and private loop devices on the additional Arbitrated Loop.
3. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system is a Fiber
30 Channel fabric.
4. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system is an intelligent bridging hub.

5. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system implements route filtering.

5

6. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system implements port ALPA range configuration.

7. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system implements I/O probe spoofing.

10

8. The interconnection system of claim 1 or 2 wherein the Fibre Channel private loop device interconnect system resets a port which has a host connected to it upon the addition of storage on an Arbitrated Loop connected to an additional Fibre Channel private loop device interconnect system port.

15

9. The interconnection system of claim 1 or 2 where there are 2 loops.

20

10. The interconnection system of claim 1 or 2 where there are more than 10 loops.

11. The interconnection system of claim 1 or 2 wherein the private loop devices include storage.

25

12. The interconnection system of claim 11 wherein the storage is JBODs.

30

13. The interconnection system of claim 11 wherein the storage is RAIDs.

14. The interconnection system of claim 11 wherein the storage is tape drives.

5 15. The interconnection system of claim 1 or 2 wherein the private loop devices include hosts.

16. The interconnection system of claim 1 or 2 further including one or more public loop devices.

10 17. The interconnection system of claim 1 or 2 wherein the private loop devices includes a bridge.

18. The interconnection system of claim 17 wherein the bridge is a SCSI to Fibre Channel bridge.

15 19. The interconnection system of claim 1 or 2 wherein the private loop devices includes a router.

20 20. The interconnection system of claim 19 wherein the router connects Fibre Channel to asynchronous transfer mode (ATM).

21. The interconnection system of claim 19 wherein the router connects Fibre Channel to ethernet.

25 22. The interconnection system of claim 21 wherein the router connects Fibre Channel to fast ethernet.

30 23. The interconnection system of claim 21 wherein the router connects Fibre Channel to a gigabit ethernet.

24. The interconnection system of claim 1 or 2 wherein at least one arbitrated loop connection includes a hub.

25. The interconnection system of claims 1 or 2 wherein there are redundant connections between a Fibre Channel private loop device interconnect system and an attached storage unit or host unit.

5

26. An interconnection system for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

10 first Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop, and
a second Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop, the first Fibre Channel private loop device interconnect system and second Fibre Channel private loop device interconnect system being directly connected, for routing frames between
15 the private loop device on the first Fibre Channel private loop device interconnect system with a private loop device on the second Fibre Channel private loop device interconnect system.

20

27. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

25 a first Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop,
a second Fibre Channel private loop device interconnect system, and
a third Fibre Channel private loop device interconnect system, with at least one Arbitrated Loop,
30 the first Fibre Channel private loop device interconnect system being directly connected to the second Fibre Channel private loop device interconnect system, which is directly connected to the third Fibre Channel private loop device interconnect system, for routing frames

between the private loop devices on the first Fibre Channel private loop device interconnect system with a private loop device on the second Fibre Channel private loop device interconnect system.

5

28. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

10 a first Fibre Channel private loop device interconnect system, with M ports,

M -1 second Fibre Channel private loop device interconnect system, the first Fibre Channel private loop device interconnect system directly connecting to each of the M-1 Fibre Channel private loop device interconnect system, and

15

a JBOD connected to the first Fibre Channel private loop device interconnect system.

29. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

20

a first Fibre Channel private loop device interconnect system with M ports, and

25 M second Fibre Channel private loop device interconnect system, each second Fibre Channel private loop device interconnect system connected to the first Fibre Channel private loop device interconnect system for routing frames between a private loop device on one of the second Fibre Channel private loop device interconnect systems and

30

a private loop device on another of the second Fibre Channel private loop device interconnect systems through the first Fibre Channel private loop device interconnect system.

30. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first Fibre Channel private loop device interconnect system,

a second Fibre Channel private loop device interconnect system, and

at least a first private loop device on a first Fibre Channel Arbitrated Loop, the first device connected to the first Fibre Channel private loop device interconnect system and to the second Fibre Channel private loop device interconnect system, and

at least a second private loop device on a second Fibre Channel arbitrated loop, the first Fibre Channel arbitrated loop and second Fibre Channel arbitrated loop being physically separate from each other, the second device being connected to the first Fibre Channel private loop device interconnect system and the second Fibre Channel private loop device interconnect system, for routing frames between the first device on the first Fibre Channel arbitrated loop and the second device on the second Fibre Channel arbitrated loop.

31. An Interconnection System for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising:

a first set of Fibre Channel private loop device interconnect systems,

a second set of Fibre Channel private loop device interconnect systems,

a first intermediate Fibre Channel private loop device interconnect system, and

a second intermediate Fibre Channel private loop device interconnect system,

each of the first set of Fibre Channel private loop device interconnect systems connected to the first intermediate Fibre Channel private loop device interconnect system, and separately to the second intermediate Fibre Channel private loop device interconnect system, and

each of the second set of Fibre Channel private loop device interconnect systems connected to the first intermediate Fibre Channel private loop device interconnect system, and separately to the second intermediate Fibre Channel private loop device interconnect system for routing frames between private loop devices in the first set and private loop devices in the second set.

32. A method for connecting a plurality of physically separate Fibre Channel Arbitrated Loops, the loops containing one or more private loop devices, comprising the steps of:

providing a first Arbitrated Loop containing one or more private loop devices,

providing one or more Arbitrated Loops containing one or more private loop devices,

disposing a Fibre Channel private loop device interconnect system between the first Arbitrated Loop and an additional Arbitrated Loop, and

routing frames between private loop devices on the First Arbitrated Loop and private loop devices on the Second Arbitrated Loop.

33. A method for implementing a logical loop of private loop devices, comprising the steps of:

segmenting the logical loop of private loop devices into a plurality of sets,

assigning each set to a physical Arbitrated Loop,

connecting the Arbitrated Loops to a Fibre Channel private loop device interconnect system to interconnect the Arbitrated Loops.

5 34. A method for routing frames between a private loop device on a first Arbitrated Loop and one or more private loop devices on at least a second Arbitrated Loop comprising the steps of:

10 receiving frames over the first Arbitrated Loop at a connected port of a Fiber Channel private loop device interconnect system, and

 filtering said frames by:

15 (1) forwarding the frame on the first Arbitrated Loop if it has an address on the first Arbitrated Loop and

 (2) providing an "open" response on the first Arbitrated Loop, if the address is not on the first Arbitrated Loop.

20 35. The method for routing frames of claim 34 further including the step of attempting to route the frame to the address which is not on the first Arbitrated Loop.

25 36. The method for routing frames of claim 34 further including the step of buffering of the frames destined to private loop devices not on the first Arbitrated Loop if the route cannot be made immediately.

30 37. The method for routing frames of claim 34 further including the step of look at destination address and compare with predefined table, then cut through if the route can be made immediately.

38. The method of claim 34 wherein the filtering is achieved by:

applying a mask to the received ALPA,

comparing the results with preconfigured value, and

5 comparing the result (match/no match) with a filter on the match/no match bit.

39. A method for configuring of multiple Fibre Channel Arbitrated Loops of private loop devices, each private loop device on an Arbitrated Loop having an Arbitrated Loop Physical Address (ALPA), comprising the steps of:

dividing the ALPAs into non-overlapping sets,

assigning each set to a separate physical Arbitrated Loop, and

15 during loop initialization, forcing the attached private loop devices to choose from the assigned set.

40. The method of claim 39 further including the step of interconnecting each subplurality through a Fibre Channel private loop device interconnect system.

41. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop devices, the loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

25 detecting an addition of a storage device to a first Arbitrated Loop, and

30 resetting the Arbitrated Loops on which hosts reside upon such detected addition.

42. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop devices, the
5 loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

receiving PLOGI I/O probes at a Fibre Channel private loop device interconnect system,

performing address lookup for the received PLOGI I/O
10 probes, and if a match exists in the lookup,

routing I/O probes from Fibre Channel SCSI initiator to private loop devices on the Fibre Channel private loop device interconnect system or other Fibre Channel private loop device interconnect systems.
15

43. A method for operation of an interconnection system, the system including more than one Arbitrated Loop, at least one loop adapted to contain storage and one loop containing a host, the devices attached to the loops being private loop devices, the
20 loops being connected to a Fibre Channel private loop device interconnect system, comprising the steps of:

receiving PLOGI I/O probes at a Fibre Channel private loop device interconnect system,

performing address lookup for the received PLOGI I/O
25 probes, and

routing the I/O probes to the Fibre Channel private loop device interconnect system controller and returning a link service reject if either there is no match found upon address lookup or if a match is found upon address lookup,
30 but no device with the destination ALPA exists on the Arbitrated Loop corresponding to the destination.

44. The method of claim 43 wherein the determination that no device with the destination ALPA exists on the Arbitrated Loop is performed by the steps of:

5 sending a class 3 frame on a loop, and
 simultaneously receiving the same frame on the same loop.

45. A bridging hub for connection of a plurality of Arbitrated Loops, the Arbitrated Loops being adapted to contain private loop devices, comprising:

10 at least first and second hub submodules,
 the submodules comprising:
 a plurality of ports, the ports including
 port bypass circuits connected to the ports
15 for connection to the Arbitrated Loops
 adapted to contain private loop devices,
 and
 an ALPA filtering port,
 a router, the router disposed between the first and
20 second hub submodules, and
 a processor control coupled to the router and the
 first and second submodules.

46. The bridging hub of claim 45 wherein the router does
25 not support class 1 connections.

47. The bridging hub of claim 45 wherein the processor control does not provide backup route determination mechanisms.

STATEMENT UNDER ARTICLE 19

The amendments made in the accompanying Submission Under Article 19 are fully supported in the specification as originally filed. For example, the specification at page 13, lines 22-24 defines a "FC-PLD-IS" or Fibre Channel Private Loop Device Interconnect System as "a fabric or intelligent bridging hub, or a device which achieves the functionality of these devices". A "fabric" is in turn defined, at least in part, as being capable of "routing frames by using the destination identifier (D_ID) information in the FC-2 frame header". (page 13, lines 15-18). Similarly, the definition of "intelligent bridging hub" at page 13, lines 19-21 also includes the function, in part, of "routing functions".

1/16

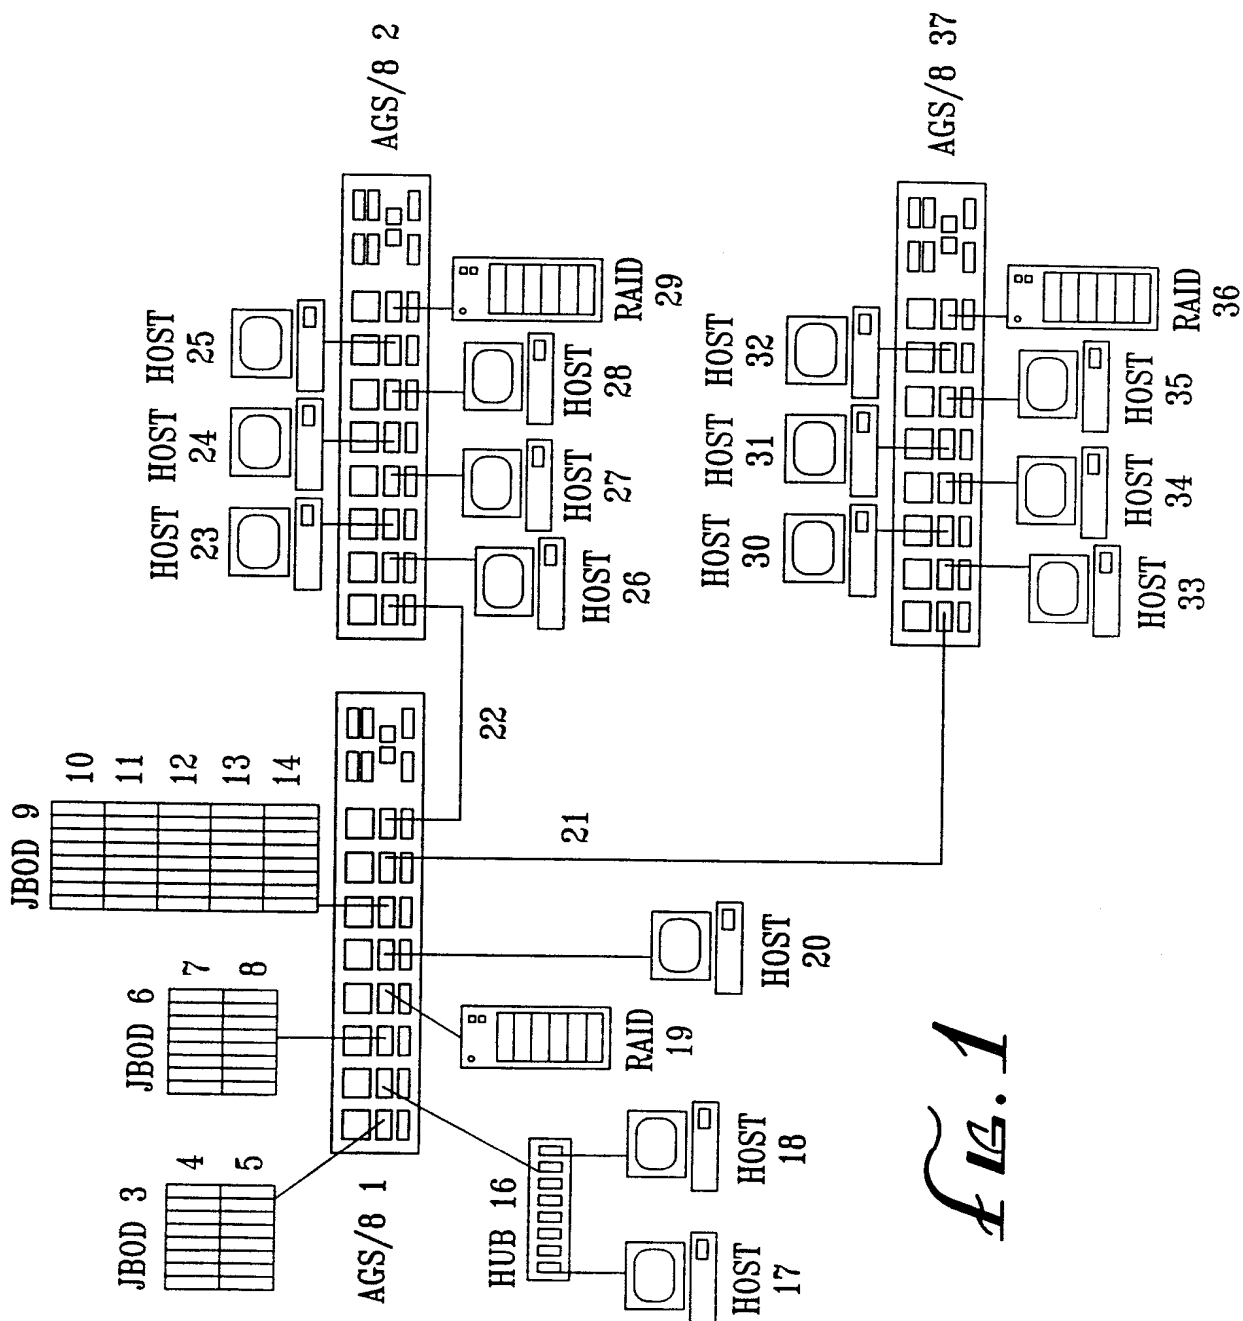


Fig. 1


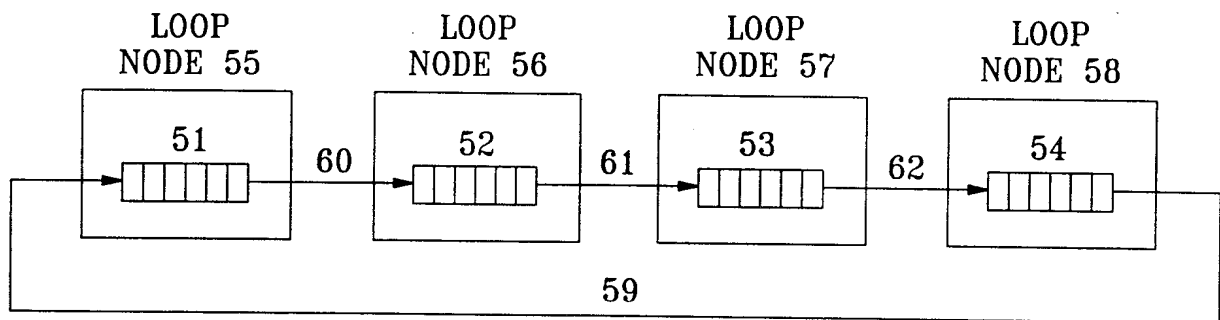
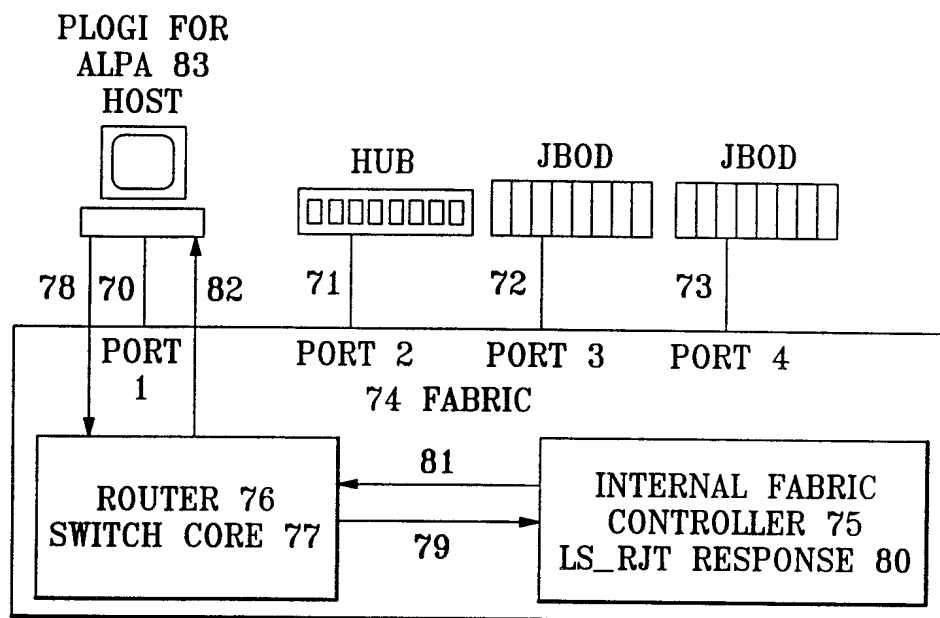




Fig. 2


2/16

*FIG. 3**FIG. 4*

3/16

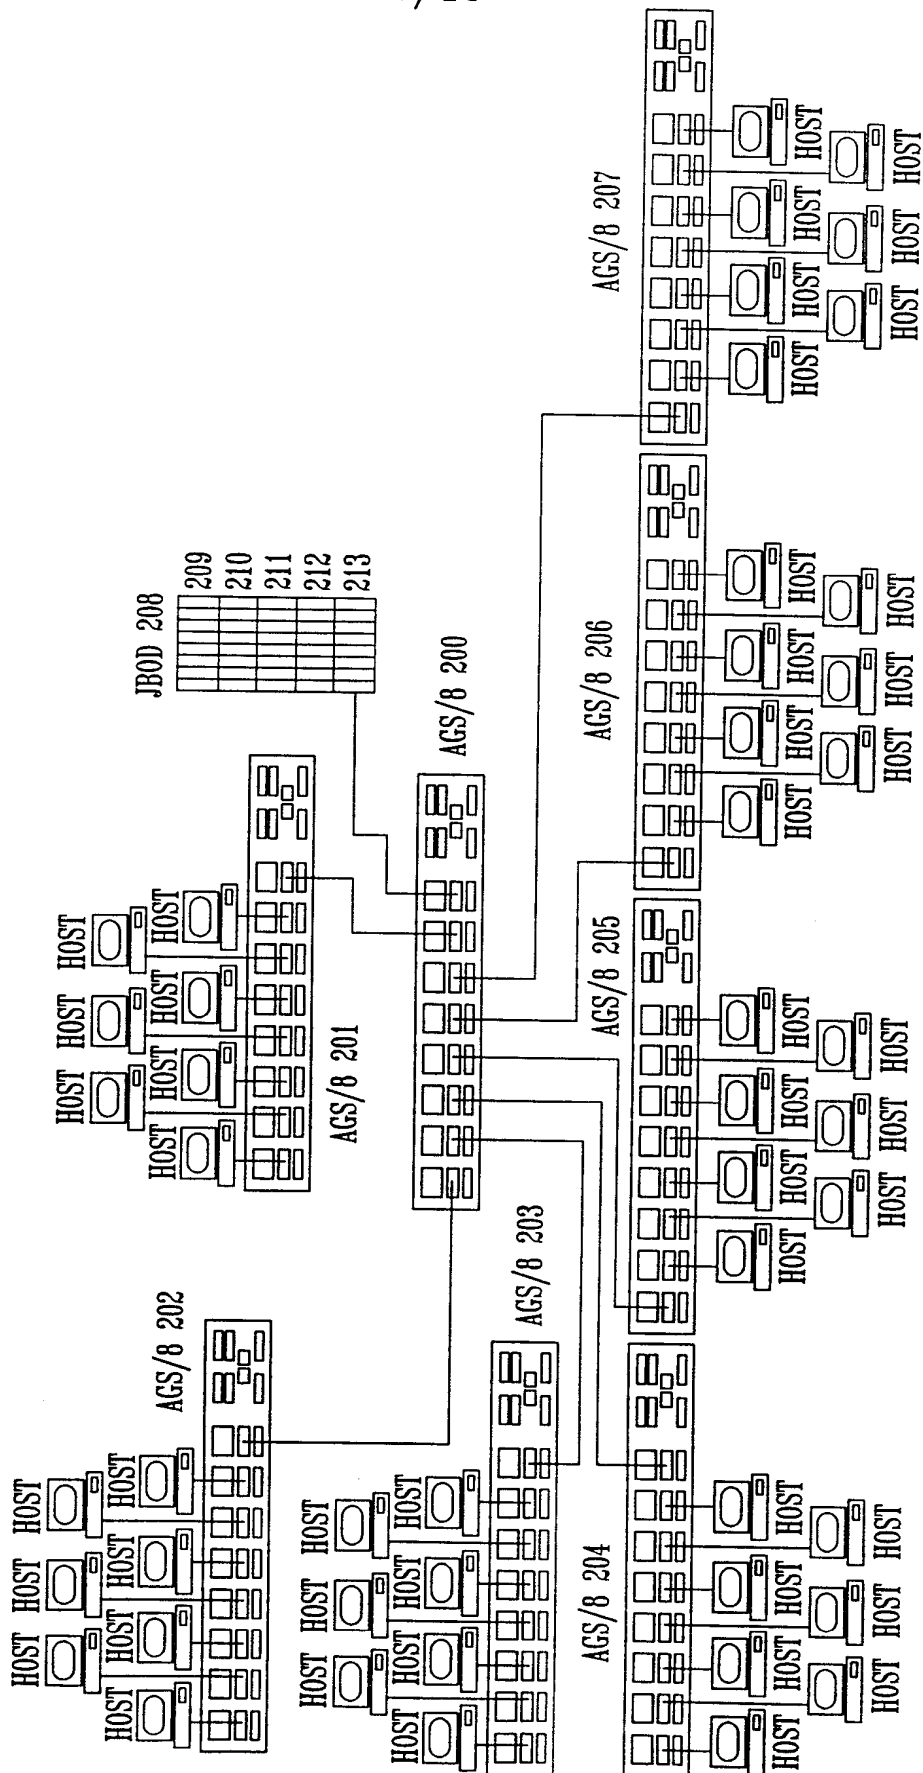
*Fig. 5**Fig. 6*

4/16



5/16

Fig. 8



6/16




Fig. 9




Fig. 11

7/16

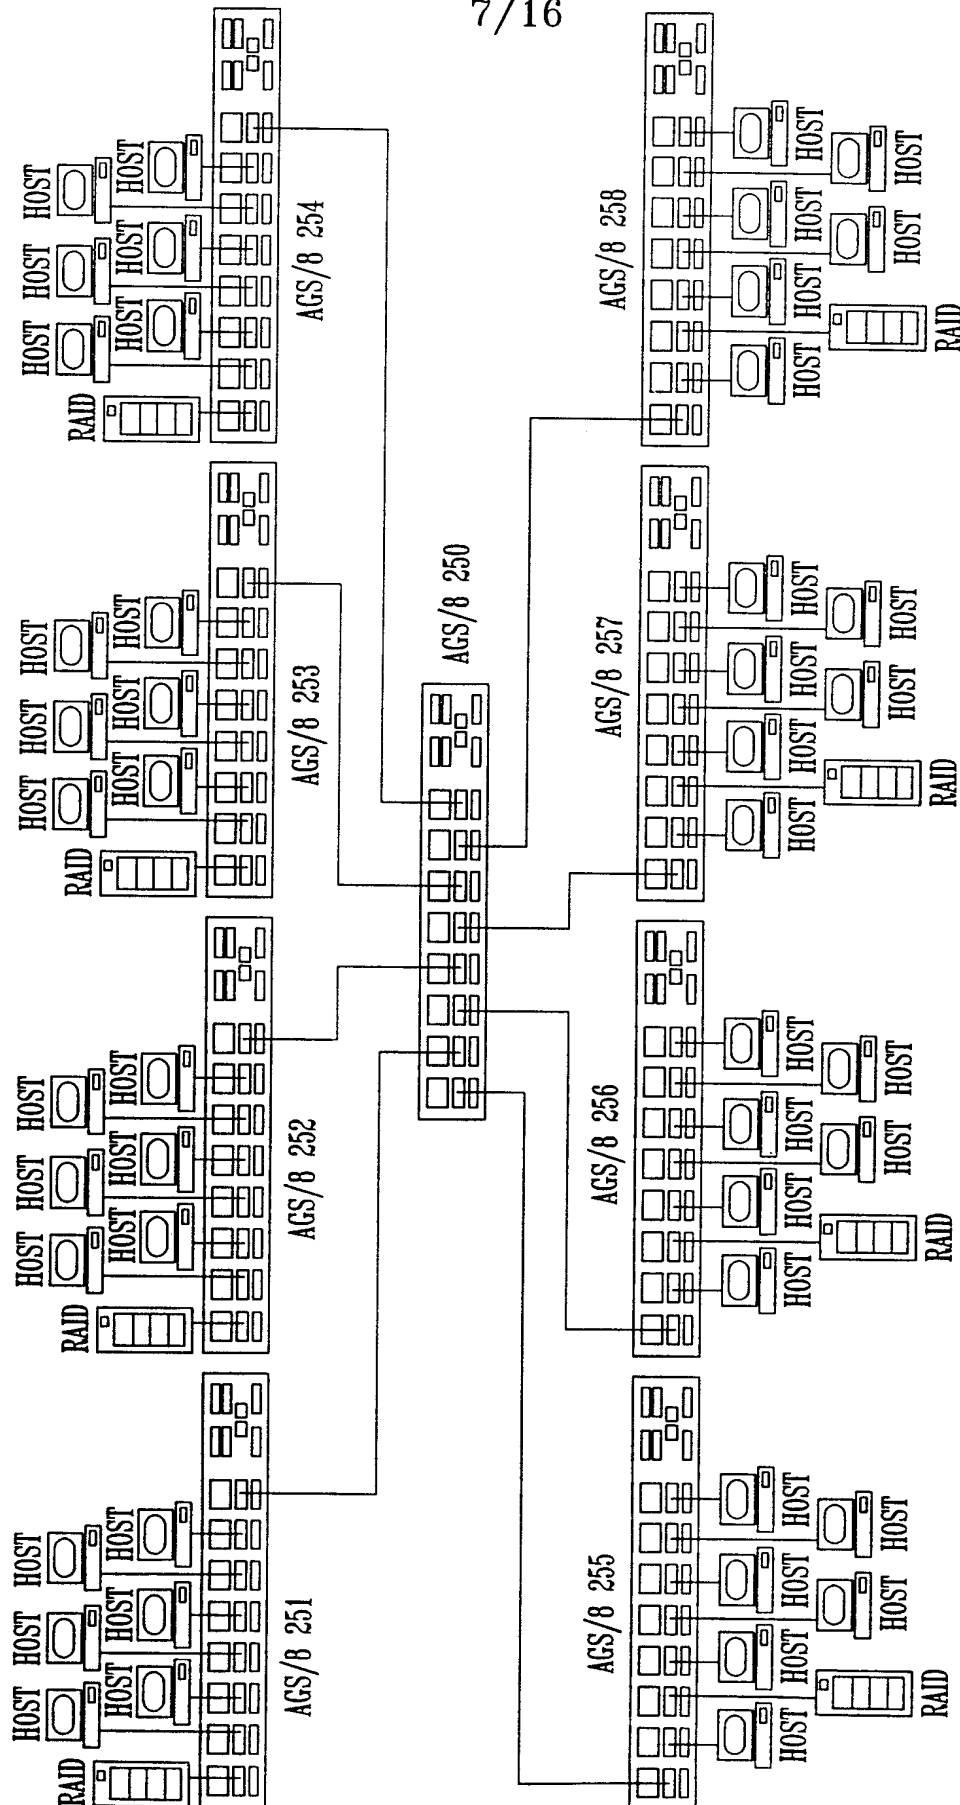


Fig. 10

8/16

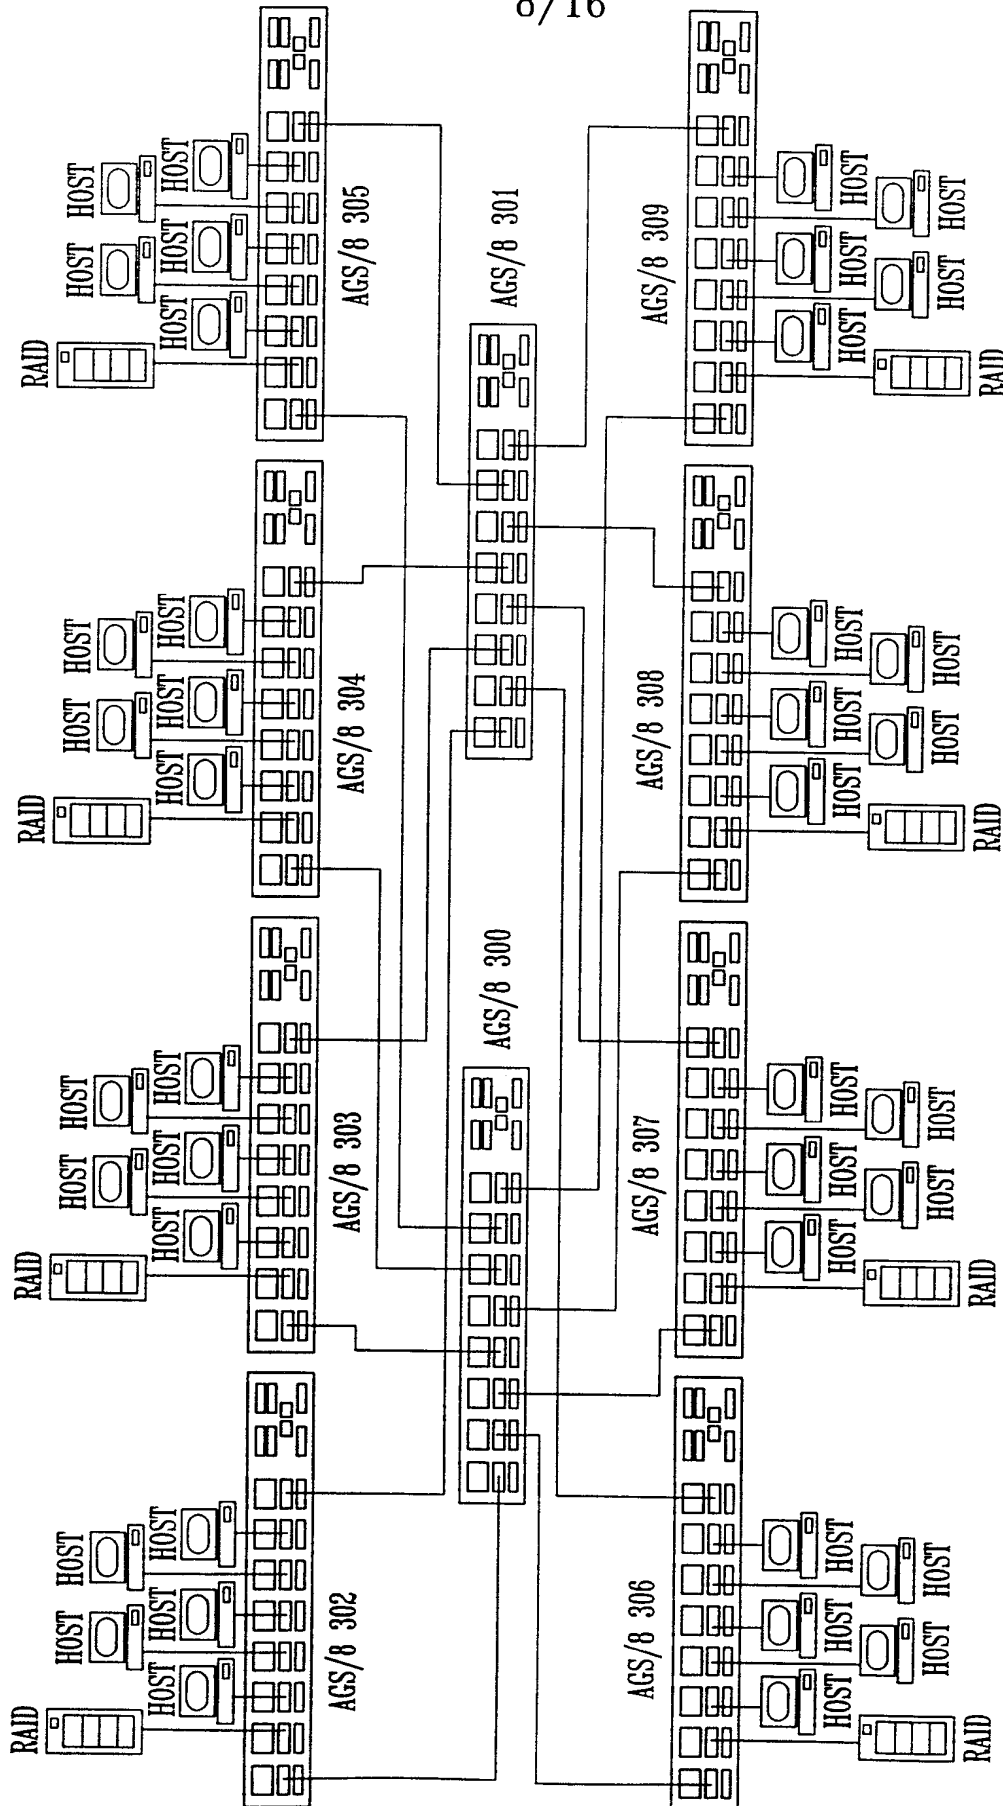

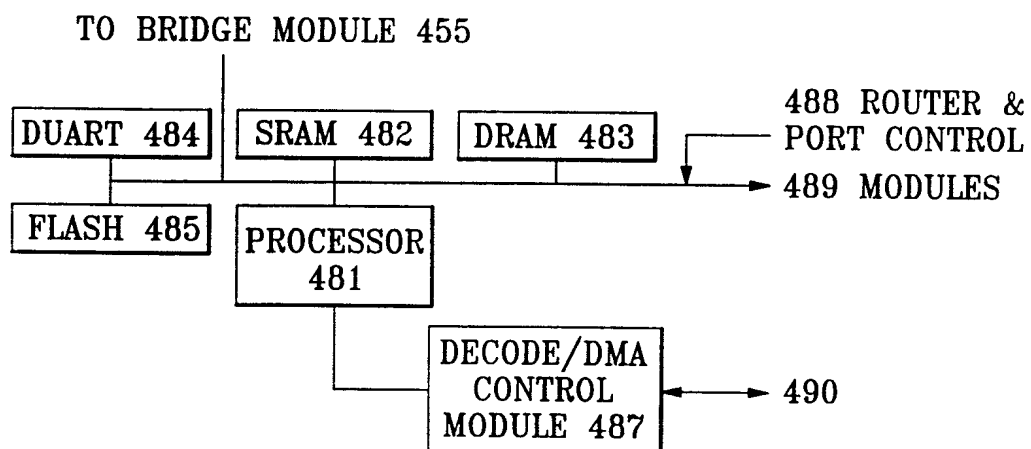


Fig. 12


9/16


$$f_{13}$$

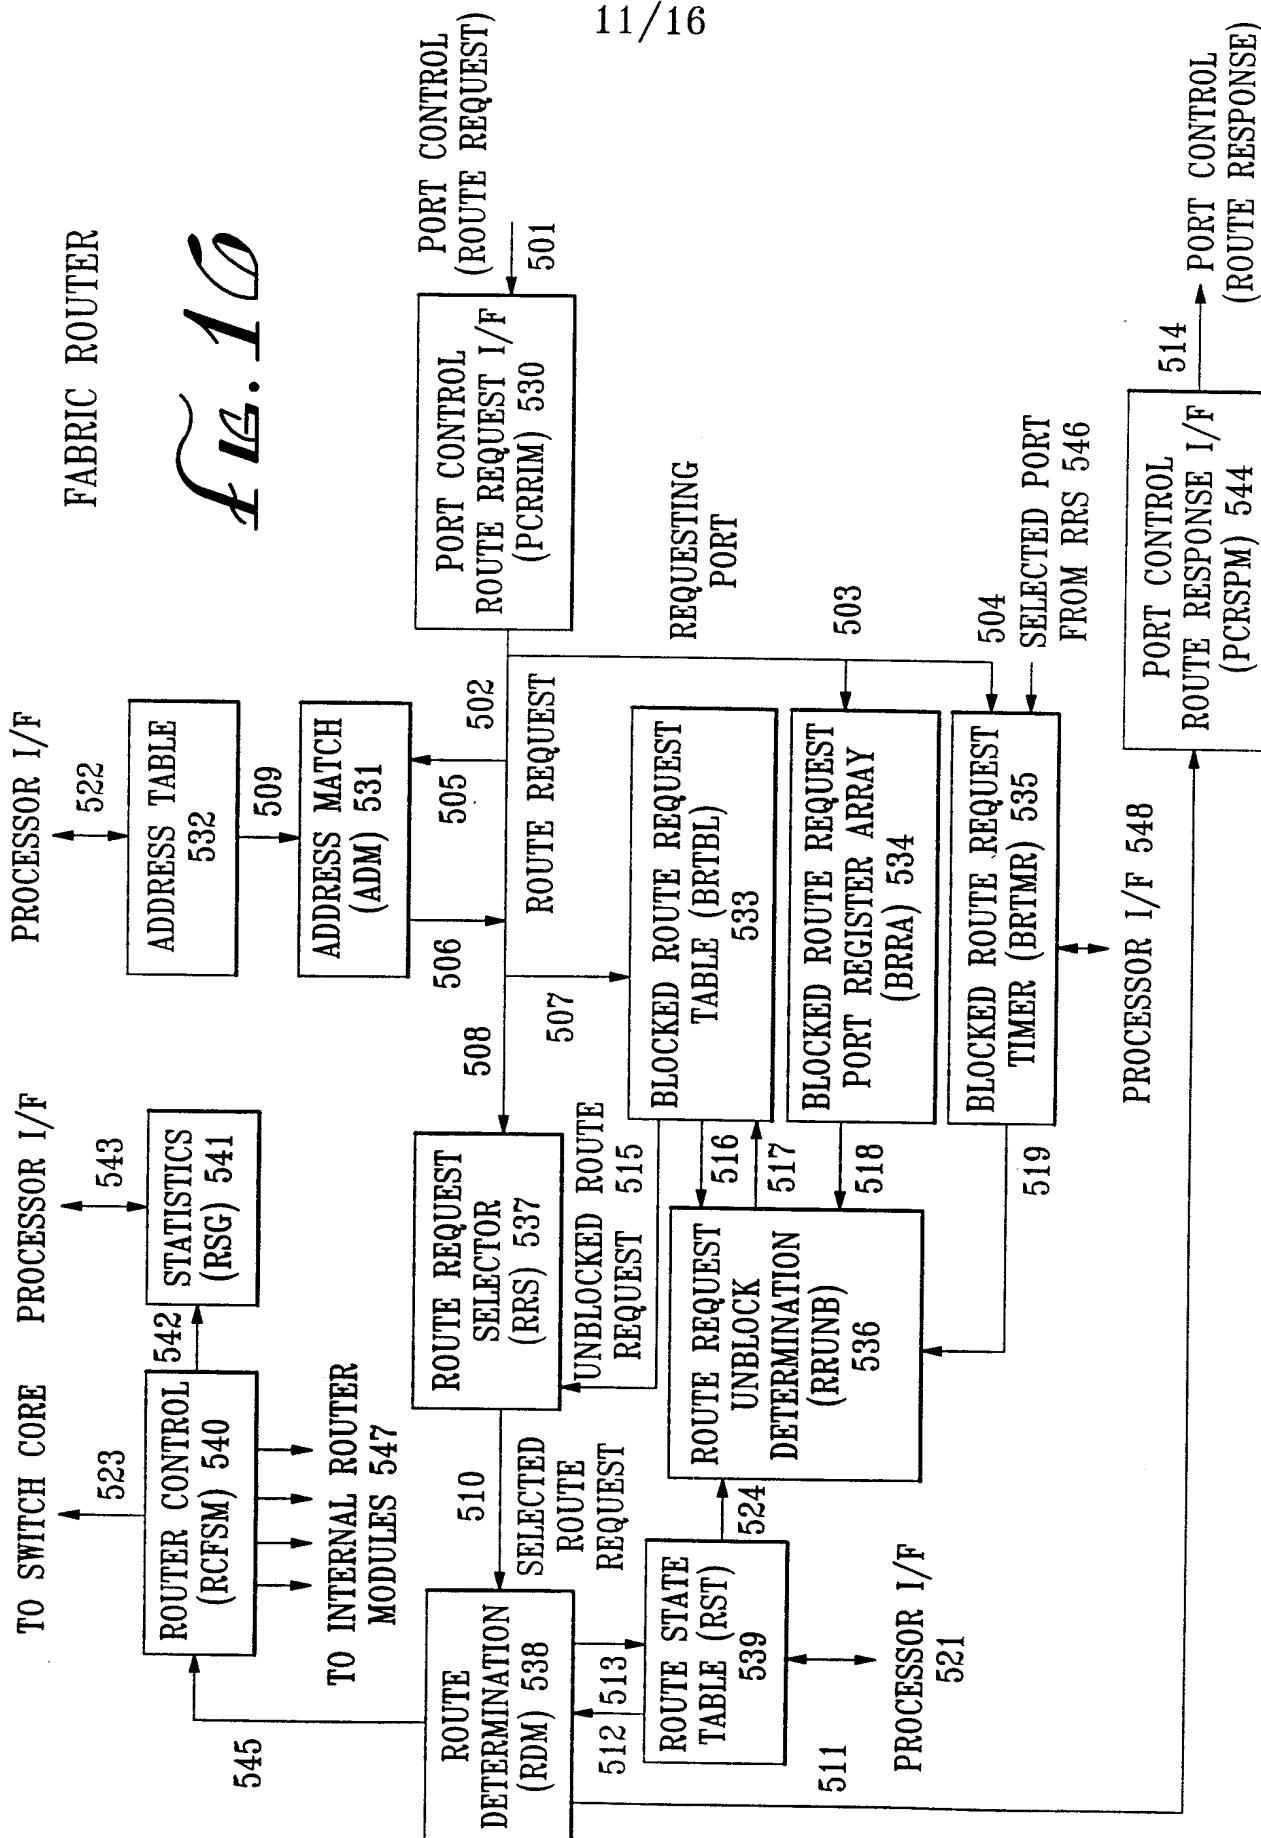
FABRIC CONTROL


$$f_{14.15}$$


10/16




11/16



12/16



LIFA FRAME FORMAT



LIFA FRAME FORMAT

EXAMPLE ALLOWING ONLY THE ALPA'S 10, 17, 18, 1B, 1D, 1E AND 1F TO BE CHOSEN





FIG. 19

LISM FRAME FORMAT

EXAMPLE WORLD WIDE PORT NAME OF ALL ZEROS TO OBTAIN LOOP INITIALIZATION MASTERSHIP



610

FIG. 20

14/16

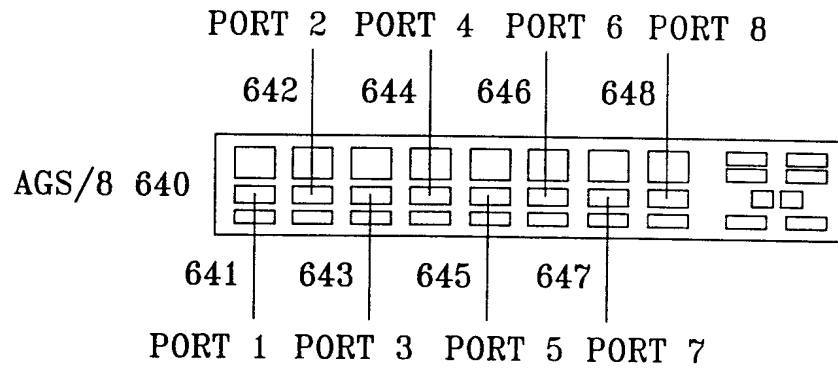


Fig. 21






Fig. 24

15/16

*FIG. 22**FIG. 23*




fig. 25

INTERNATIONAL SEARCH REPORT

Int ional Application No
PCT/US 98/15554

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L12/46 H04L12/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	G.R. STEPHENS ET AL.: "Fibre Channel. The basics" June 1995, ANCOT CORPORATION, USA XP002085982	1-3, 26, 29, 30, 32, 41
Y	see page 5-6 - page 5-9	4, 5, 11-14, 17-21, 28 10, 15, 27
A	---	
Y	RICKARD W: "FIBRE CHANNEL AS A NETWORK BACKBONE" WESCON '94, 27 September 1994, pages 653-659, XP000532640	5, 11-14, 17-21, 28
A	see paragraph 3.1	8, 22, 23, 42
	see paragraph 3.2	

	-/--	

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

27 November 1998

Date of mailing of the international search report

09/12/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

Int ional Application No
PCT/US 98/15554

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	CHIN H: "FIBRE CHANNEL OFFERS ANOTHER ROAD TO HIGH-SPEED NETWORKING AND I/O" COMPUTER TECHNOLOGY REVIEW, vol. 15, no. SUPPL, 21 December 1995, pages 45-47, XP000502896	4
A	see the whole document	24, 45
A	ANUJAN VARMA ET AL: "USING CAMP-ON TO IMPROVE TE PERFORMANCE OF A FIBRE CHANNEL SWITCH" PROCEEDINGS OF THE CONFERENCE ON LOCAL COMPUTER NETWORKS, 1 January 1993, pages 247-255, XP000562898 see page 247, right-hand column, line 3 - line 27 see paragraph 2.4 see paragraph 3	34-38, 43
A	MARTIN C R: "FABRIC INTERCONNECTION OF FIBRE CHANNEL STANDARD NODES" PROCEEDINGS OF THE SPIE, vol. 1784, 8 September 1992, pages 65-71, XP000562875 see page 68, line 27 - line 42 see paragraph 6	33, 39
A	RAVINDRAN G ET AL: "A COMPARISON OF BLOCKING AND NON-BLOCKING PACKET SWITCHING TECHNIQUES IN HIERARCHICAL RING NETWORKS" IEICE TRANSACTIONS ON INFORMATION AND SYSTEMS1, vol. E79-D, no. 8, August 1996, pages 1130-1138, XP000628514 see page 1130, right-hand column, line 6 - line 27 see paragraph 2.1 see figure 1	27, 31, 37